

UNIVERSIDAD AUTÓNOMA DE MADRID
ESCUELA POLITÉCNICA SUPERIOR

ANALYSIS AND STUDY ON TEXT
REPRESENTATION TO IMPROVE
THE ACCURACY OF THE
NORMALIZED COMPRESSION
DISTANCE

THESIS

Ana Granados Fontecha
Madrid, 2012

ADVISORS: David Camacho Fernández
Francisco de Borja Rodríguez Ortiz

Tribunal nombrado por el Mgfco. y Excmo. Sr. Rector de la Universidad
Autónoma de Madrid, el día de de 2012.

Presidente: D.

Vocal: D.

Vocal: D.

Vocal: D.

Secretario: D.

Realizado el acto de defensa y lectura de la Tesis el día de
de 2012 en

Calificación:

EL PRESIDENTE

EL SECRETARIO

LOS VOCALES

A mis padres.

*After climbing a great hill, one only finds
that there are many more hills to climb.*

-Nelson Mandela

Abstract

The huge amount of information stored in text form makes methods that deal with texts really interesting. This thesis focuses on dealing with texts using compression distances. More specifically, the thesis takes a small step towards understanding both the nature of texts and the nature of compression distances. Broadly speaking, the way in which this is done is exploring the effects that several distortion techniques have on one of the most successful distances in the family of compression distances, the Normalized Compression Distance -NCD-.

The research carried out in this thesis can be divided into three parts. The first part, which corresponds to Chapter 5, experimentally evaluates the impact that several word removal techniques have on NCD-driven text clustering, with the aim of better understanding of both the nature of compression distances and the nature of textual information. This goal is accomplished by analyzing how the information contained in the documents and how the upper bound estimation of their Kolmogorov complexity progress as words are removed from the documents. One of the main conclusions that can be drawn from this analysis is that the clustering accuracy can be improved by applying a specific word removal technique. This distortion technique consists of removing the most frequent words of the language preserving the previous text structure.

The second part of the thesis, which corresponds to Chapter 6, attempts to shed light on the reasons why the application of such a distortion technique can improve NCD-driven text clustering. The experimental results show that the maintenance of both the previous text structure and the remaining words structure have some relevance in the clustering behavior.

The third part of the thesis, which corresponds to Chapter 7, applies the above mentioned distortion technique to NCD-driven document search. The application of compression distances to document search is not trivial due to the fact that they do not commonly perform well when the compared objects have very different sizes. An NCD-based document search engine that deals with that drawback by using passage retrieval, is used in the third part of the thesis. The results show that the search accuracy can be improved by applying the distortion technique presented previously.

Summarizing, one of the distortion techniques explored in the thesis has been found to be beneficial both in NCD-based document clustering and in NCD-based document search.

Contents

0	Resumen de la Tesis	1
1	Introduction	7
2	Objectives	13
3	Thesis Overview	15
4	Related Work	17
4.1	Information Theory Concepts	18
4.1.1	Kolmogorov Complexity	22
4.2	Compression Algorithms	23
4.2.1	Statistical Methods	24
4.2.2	Dictionary Methods	29
4.2.3	Other Methods	32
4.2.4	Comparing Compressors: Calgary Corpus	36
4.3	Compression Distances	36
4.3.1	Analyzing some extreme cases	39
4.3.2	Understanding NCD	40
4.3.3	Some NCD applications	42
4.4	Text Distortion Techniques	45
4.5	Contextual Information	47
5	Study on text distortion	49
5.1	Distortion Techniques	50
5.2	Experimental Setup	57
5.2.1	NCD-based Text Clustering	57
5.2.2	Datasets	60
5.3	Experimental Results	60
5.3.1	The Books dataset and the PPMZ compressor	61
5.3.2	Results for the <i>asterisk substitution method</i>	67

5.3.3	Synopsis of all the obtained results	70
5.4	Summary and Conclusions	72
6	Relevance of contextual information	75
6.1	Distortion Techniques	76
6.2	Experimental Setup	78
6.2.1	Datasets	78
6.3	Experimental Results	78
6.3.1	Synopsis of results	83
6.4	Summary and Conclusions	87
7	Application to Document Searching	89
7.1	NCD-based Document Search Method	90
7.2	Datasets	91
7.3	Experimental Results	92
7.3.1	Summary of Results	93
7.4	Summary and Conclusions	98
8	Conclusions	99
9	Summary of Results	101
10	Contributions	105
A	Acronyms	109
B	Datasets	111
B.1	Books dataset	111
B.2	UCI-KDD dataset	112
B.3	MedlinePlus dataset	113
B.4	IMDB dataset	114
B.5	SRT-serial dataset	115
B.6	UCM dataset	117
B.7	Reuters dataset	118
B.8	20newsgroups dataset	119
C	Queries	123
D	Detailed Experimental Results	127
D.1	Preliminary study on text distortion	127

List of Figures

4.1	Huffman tree generation.	25
4.2	PPM: seven tries of “bananas” for a context of $N = 2$	28
5.1	Visual representation of the information loss.	56
5.2	Percentage of substituted words.	57
5.3	Example of dendrogram for the Books repository.	58
5.4	Estimation of an upper bound for the Books complexity.	61
5.5	Books. PPMZ compressor. MFW selection method.	62
5.6	Books. PPMZ compressor. RW selection method.	63
5.7	Books. PPMZ compressor. LFW selection method.	64
5.8	Dendrogram obtained with no distortion.	65
5.9	Perfect dendrogram.	66
5.10	Understanding the tables.	69
6.1	Text distortion techniques.	77
6.2	Clustering results for the UCI-KDD dataset.	79
6.3	Clustering results for the Books dataset.	80
6.4	Clustering results for the MedlinePlus dataset.	81
6.5	Clustering results for the IMDB dataset.	82
6.6	Clustering results for the SRT-serial dataset.	83
6.7	Clustering error difference with the <i>Original sorting</i> distortion technique.	84
7.1	UCM dataset. Benefits of applying distortion.	94
7.2	20newsgroups dataset. Benefits of applying distortion.	95
7.3	Reuters dataset. Benefits of applying distortion.	96
7.4	Percentage of improvement.	97
B.1	Books. Extract from <i>Don Quixote</i> by Miguel de Cervantes.	112
B.2	UCI-KDD. Extract from a document on cryptography.	113
B.3	MedlinePlus. Extract from a document on diabetes.	114
B.4	IMDB. Extract from the movie <i>The Matrix</i>	116

B.5	SRT-serial. Extract from the serial <i>Heroes</i>	117
B.6	UCM. Extract from a paper on grid computing.	118
B.7	Reuters. Extract from a document on wheat.	120
B.8	20newsgroups. Extract from a document on <i>sci.space</i>	121
C.1	UCM. Example of query.	125
C.2	20newsgroups. Example of query.	125
C.3	Reuters. Example of query.	126
D.1	Books. PPMZ compressor. MFW selection method.	128
D.2	Books. PPMZ compressor. RW selection method.	128
D.3	Books. PPMZ compressor. LFW selection method.	129
D.4	Books. PPMZ compressor. Dendrogram obtained with no distortion.	130
D.5	Books. PPMZ compressor. Best dendrogram obtained.	131
D.6	Books. LZMA compressor. MFW selection method.	132
D.7	Books. LZMA compressor. RW selection method.	132
D.8	Books. LZMA compressor. LFW selection method.	133
D.9	Books. LZMA compressor. Dendrogram obtained with no distortion.	134
D.10	Books. LZMA compressor. Best dendrogram obtained.	135
D.11	Books. BZIP2 compressor. MFW selection method.	136
D.12	Books. BZIP2 compressor. RW selection method.	136
D.13	Books. BZIP2 compressor. LFW selection method.	137
D.14	Books. BZIP2 compressor. Dendrogram obtained with no distortion.	138
D.15	Books. BZIP2 compressor. Best dendrogram obtained.	139
D.16	UCI-KDD. PPMZ compressor. MFW selection method.	140
D.17	UCI-KDD. PPMZ compressor. RW selection method.	140
D.18	UCI-KDD. PPMZ compressor. LFW selection method.	141
D.19	UCI-KDD. PPMZ compressor. Best dendrogram obtained.	142
D.20	UCI-KDD. LZMA compressor. MFW selection method.	143
D.21	UCI-KDD. LZMA compressor. RW selection method.	143
D.22	UCI-KDD. LZMA compressor. LFW selection method.	144
D.23	UCI-KDD. LZMA compressor. Best dendrogram obtained.	145
D.24	UCI-KDD. BZIP2 compressor. MFW selection method.	146
D.25	UCI-KDD. BZIP2 compressor. RW selection method.	146
D.26	UCI-KDD. BZIP2 compressor. LFW selection method.	147
D.27	UCI-KDD. BZIP2 compressor. Best dendrogram obtained.	148
D.28	MedlinePlus. PPMZ compressor. MFW selection method.	149
D.29	MedlinePlus. PPMZ compressor. RW selection method.	149

D.30 MedlinePlus. PPMZ compressor. LFW selection method. . . .	150
D.31 MedlinePlus. PPMZ compressor. Dendrogram obtained with no distortion.	151
D.32 MedlinePlus. PPMZ compressor. Best dendrogram obtained. .	152
D.33 MedlinePlus. LZMA compressor. MFW selection method. . .	153
D.34 MedlinePlus. LZMA compressor. RW selection method. . . .	153
D.35 MedlinePlus. LZMA compressor. LFW selection method. . . .	154
D.36 MedlinePlus. LZMA compressor. Dendrogram obtained with no distortion.	155
D.37 MedlinePlus. LZMA compressor. Best dendrogram obtained. .	156
D.38 MedlinePlus. BZIP2 compressor. MFW selection method. . .	157
D.39 MedlinePlus. BZIP2 compressor. RW selection method. . . .	157
D.40 MedlinePlus. BZIP2 compressor. LFW selection method. . . .	158
D.41 MedlinePlus. BZIP2 compressor. Dendrogram obtained with no distortion.	159
D.42 MedlinePlus. BZIP2 compressor. Best dendrogram obtained. .	160
D.43 IMDB. PPMZ compressor. MFW selection method.	161
D.44 IMDB. PPMZ compressor. RW selection method.	161
D.45 IMDB. PPMZ compressor. LFW selection method.	162
D.46 IMDB. PPMZ compressor. Best dendrogram obtained. . . .	163
D.47 IMDB. LZMA compressor. MFW selection method.	164
D.48 IMDB. LZMA compressor. RW selection method.	164
D.49 IMDB. LZMA compressor. LFW selection method.	165
D.50 IMDB. LZMA compressor. Dendrogram obtained with no dis- tortion.	166
D.51 IMDB. LZMA compressor. Best dendrogram obtained. . . .	167
D.52 IMDB. BZIP2 compressor. MFW selection method.	168
D.53 IMDB. BZIP2 compressor. RW selection method.	168
D.54 IMDB. BZIP2 compressor. LFW selection method.	169
D.55 IMDB. BZIP2 compressor. Best dendrogram obtained. . . .	170

List of Tables

4.1	LZ77: Lempel-Ziv sliding window.	32
4.2	Move-To-Front transform.	34
4.3	Burrows-Wheeler Transform encoding.	36
4.4	Calgary Corpus.	37
4.5	Comparison of compression algorithms.	37
4.6	Understanding NCD: Text samples.	41
4.7	Understanding NCD: matrix distances.	42
4.8	Helping the compressor.	46
5.1	<i>MFW selection method & asterisk substitution method.</i>	52
5.2	<i>RW selection method & asterisk substitution method.</i>	53
5.3	<i>LFW selection method & asterisk substitution method.</i>	54
5.4	Clustering error measurement.	59
5.5	MFW selection method	68
5.6	RW selection method	68
5.7	LFW selection method	68
5.8	Books dataset. Average clustering error	71
5.9	UCI-KDD dataset. Average clustering error	71
5.10	MedlinePlus dataset. Average clustering error	71
5.11	IMDB dataset. Average clustering error	71
6.1	Average clustering error.	86
6.2	Normalized average error.	86
7.1	Datasets and experiments description.	92

Chapter 0

Resumen de la Tesis

Hoy en día, la mayoría de la información almacenada electrónicamente, está almacenada en forma de texto. De hecho, si reflexionamos sobre la cantidad de tiempo que cada día pasamos delante del ordenador leyendo e-mails, noticias, artículos o informes, nos daremos cuenta que, de hecho, la mayoría de la información con la que trabajamos diariamente es texto. Esta circunstancia hace que las áreas de investigación que estudian diferentes aspectos relacionados con los datos textuales tengan más importancia cada día.

Esta tesis se centra en el tratamiento de textos mediante el uso de distancias basadas en compresión. Más específicamente, la tesis pretende avanzar en la comprensión tanto de la naturaleza de la información textual, como de las métricas basadas en compresión.

El fundamento teórico de las distancias basadas en compresión es la complejidad de Kolmogorov [66], la cual está íntimamente relacionada con el concepto de entropía propuesto por Shannon en el paper que dio lugar al nacimiento de la teoría de la información [113].

En términos generales, la teoría desarrollada por Shannon cuantifica la cantidad de información como la cantidad de sorpresa que la información contiene al ser revelada. Una forma muy simple de entender esta idea es pensar en la comunicación entre personas.

Por ejemplo, si una persona le dice a otra algo que la última ya sabía, no habrá ninguna sorpresa en el mensaje, y por tanto, la primera persona no habrá dado ninguna información a la segunda. Por el contrario, si la primera persona le dice a la segunda algo que ésta última no sabía, la primera persona le habrá dado a la última algo de información.

Ahora bien, desde el punto de vista cuantitativo, la cantidad de información transmitida en el segundo ejemplo, dependerá de lo probable que fuera el mensaje transmitido. No es lo mismo decir "Acabo de asomarme a la ventana de mi casa y he visto pasar a una persona por la calle", que

decir, "Acabo de asomarme a la ventana de mi casa y he visto pasar a la Reina de Inglaterra por la calle". De esa manera, la información definida por Shannon es inversamente proporcional a la probabilidad, es decir, cuanto menos probable sea un mensaje, más información contendrá dicho mensaje.

Para cuantificar de manera formal la información asociada a un sistema, Shannon definió el concepto de entropía como el promedio de la ganancia de información de todos los eventos posibles del sistema. Como cada evento puede ocurrir o no, con una cierta probabilidad, la entropía creada por Shannon da un peso a la información asociada a cada evento, en función a la probabilidad de dicho evento.

El concepto de entropía se ha aplicado en numerosas áreas de investigación desde su creación. En particular, la entropía es un concepto básico en el área de la compresión de datos, ya que proporciona un umbral teórico de la cantidad de compresión que se puede alcanzar al comprimir una cadena [7, 94, 103]. Este umbral teórico coincide, aproximadamente, no sólo con la entropía de la cadena, sino también con la complejidad de Kolmogorov de dicha cadena [123]. Por tanto, ambos conceptos están directamente relacionados.

La complejidad de Kolmogorov de una cadena, se define como la longitud del programa más corto que puede generar la cadena en una máquina universal de Turing [66, 76, 124]. Una cadena será más o menos compleja dependiendo de la naturaleza de la misma. Por ejemplo, la cadena "0000000000000000" será menos compleja que la cadena "0000111100001111", y a su vez, ésta será menos compleja que la cadena "1011011100101010".

La definición de complejidad de Kolmogorov puede extenderse para definir la complejidad condicional de Komogorov, la cual mide la complejidad de una cadena x relativa a otra cadena y . Esta medida se define como la longitud del programa más corto que puede generar la cadena x teniendo la cadena y como entrada a dicho programa.

Li *et al.* definieron una medida de similaridad entre dos cadenas, llamada *Normalized Information Distance* -NID-, combinando los conceptos de complejidad de Komogorov y de complejidad condicional de Kolmogorov [75].

Dado que la complejidad de Kolmogorov no es computable [123], la NID tampoco lo es. Sin embargo, Cilibrasi *et al.* propusieron una medida computable, llamada *Normalized Compression Distance* -NCD-, que utiliza algoritmos de compresión para estimar cotas superiores de la complejidad de Kolmogorov [30]. Puede encontrarse información detallada sobre la NCD y la NID en la Sección 4.3.

La NCD en particular, y las métricas basadas en compresión en general, se han aplicado a numerosas áreas de investigación debido a su naturaleza libre de parámetros, a su efectividad y a su facilidad de uso. Entre otras, las distancias basadas en compresión se han utilizado en áreas de investigación,

tales como el clustering de documentos [48, 49, 50, 51, 56, 121], la recuperación de documentos [52, 82], la clasificación de música [31, 46], la minería de datos [32], la seguridad de diferentes sistemas computacionales [4, 12, 131], la detección de plagios [26], la ingeniería del software [3, 5, 109], la bioinformática [44, 65, 69, 91], la química [85], la medicina [35, 107] o incluso el arte [119].

El hecho de que las distancias basadas en compresión se hayan utilizado tanto, da una idea de lo útiles que son. Sin embargo, a pesar de su amplio uso, se ha avanzado poco en la interpretación de sus resultados o en la explicación de su comportamiento. Cada vez que se lleva a cabo un trabajo analítico sobre las distancias basadas en compresión, normalmente éste se centra en la manipulación algebraica de conceptos algorítmicos de teoría de la información [30, 75, 139].

Uno de los propósitos de esta tesis es avanzar en el entendimiento de las métricas basadas en compresión, para así poder mejorar el rendimiento de este tipo de métricas. En particular, esta tesis se centra en una de las más importantes distancias basadas en compresión, la previamente mencionada NCD. El análisis llevado a cabo en esta tesis es principalmente experimental. Por tanto, la metodología de trabajo utilizada es la utilizada en ciencias experimentales. Esta metodología se basa en perturbar el sistema para observar las consecuencias que acarrea dicha perturbación en el estado del sistema.

La hipótesis de partida es que se puede modificar la información contenida en los textos, de manera que el compresor capture mejor la estructura de los mismos, y por tanto, se pueda mejorar el rendimiento de la NCD. La clave sería cambiar la representación de los textos sin perder la información relevante, de forma que esa nueva representación sea más favorable para que los compresores capturen mejor las similitudes entre los textos.

Antes de describir los experimentos realizados a lo largo de esta tesis y mostrar los correspondientes resultados, el Capítulo 4 presenta todos los conceptos necesarios para comprender los contenidos de la tesis. Tras la presentación de dichos conceptos, los Capítulos 5 a 7 describen los experimentos realizados a lo largo de la tesis, y muestran los resultados experimentales obtenidos. Cada uno de esos capítulos tiene un objetivo claro marcado, y genera una serie de contribuciones, las cuales se detallan siempre al comienzo de cada capítulo.

La investigación correspondiente al Capítulo 5 pretende avanzar en el entendimiento tanto de la naturaleza de la información textual, como de la naturaleza de las distancias basadas en compresión. Este avance se realiza evaluando el impacto que tienen varias técnicas de distorsión basadas en la eliminación de palabras sobre el rendimiento de la NCD.

En concreto, la investigación realizada tanto en el Capítulo 5, como en

el Capítulo 6 utiliza el método de clustering basado en la NCD desarrollado por los creadores de la NCD [29], para medir el impacto que tienen las técnicas de distorsión estudiadas sobre el rendimiento de la NCD. El uso del método de clustering basado en la NCD como herramienta para medir el rendimiento de la NCD, permite analizar cómo la información contenida en los textos estudiados evoluciona a medida que las palabras son eliminadas de los documentos.

En el Capítulo 5, además de estudiar cómo evoluciona la información contenida en los textos a medida que avanza la distorsión de los mismos, se estudia cómo la complejidad de los textos estudiados evoluciona a medida que se eliminan más y más palabras de los textos [49, 50].

Las principales contribuciones de ese capítulo pueden resumirse brevemente como sigue:

- Análisis y estudio de nuevas representaciones de datos textuales para evaluar el comportamiento de la NCD.
- Una técnica de representación de los datos textuales, especialmente diseñada para ser utilizada en herramientas que utilicen métricas basadas en compresión, que reduce la complejidad de los documentos mientras que mantiene la mayoría de la información relevante de los mismos.
- Evidencia experimental de cómo refinar la representación de los textos para permitir al compresor obtener similaridades más fiables, y por tanto, permitir al método de clustering basado en la NCD mejorar los resultados obtenidos al trabajar con los textos originales, es decir, los textos sin distorsionar.

Una de las principales conclusiones que se pueden sacar del análisis llevado a cabo en el Capítulo 5 es que la precisión del clustering se puede mejorar aplicando una de las técnicas analizadas. Esta técnica implica, no sólo la eliminación de palabras, sino también la conservación de la estructura contextual de los textos.

Esos resultados apuntan a que aunque la información más importante de un texto esté contenida en las palabras más relevantes del mismo, lo que rodea a esas palabras es importante también, ya que es el sustrato que las soporta. La hipótesis sería que esa es la razón por la cual, la técnica de distorsión que mantiene la información relevante a la vez que preserva la información contextual, es la mejor de todas las evaluadas.

El Capítulo 6 estudia si esa hipótesis es acertada o no, es decir, el capítulo estudia en qué medida se ven afectados los resultados obtenidos debido tanto

a la preservación de la información relevante, como a la preservación de la información contextual.

El concepto de información contextual se ha utilizado en numerosas aplicaciones informáticas. Por ejemplo, se ha utilizado en áreas de investigación como la recuperación de información [80, 100, 115, 117, 120], los sistemas de recomendación [2, 71, 118, 132], las aplicaciones sensibles al contexto [21, 42, 55, 96, 108], la visión artificial [1, 9, 27, 89, 90], el reconocimiento de voz [43, 58, 74, 92] o el análisis del tráfico en redes [47], entre otros.

Cuando se está trabajando con información textual, la idea de contexto es muy útil, ya que está íntimamente relacionada con los textos, debido a la naturaleza intrínseca de los mismos. Dado que los textos no son sólo secuencias de palabras, sino que tienen una estructura coherente [67], aplicar la idea de contexto al manejo de textos surge de forma natural.

En esta tesis, la información contextual es un subproducto de la aplicación de la técnica de distorsión, presentada en el Capítulo 5, mencionada anteriormente. El Capítulo 6 compara dicha técnica con tres nuevas técnicas creadas a partir de la anterior, las cuales destruyen la información contextual de diferentes maneras. Analizando los resultados experimentales obtenidos, se puede observar que mantener la información contextual es beneficioso en el campo del clustering de textos basado en la NCD [48, 51].

Las principales contribuciones del Capítulo 6 se resumen brevemente en los siguientes puntos:

- Evaluación experimental de la relevancia que la información contextual tiene en el clustering de textos basado en la NCD, en un escenario de eliminación de palabras.
- Nuevas perspectivas para la evaluación y el estudio del comportamiento de las distancias basadas en compresión, en relación a la información contextual.

Finalmente, el Capítulo 7 aplica los conocimientos adquiridos en los Capítulos 5 y 6 a la búsqueda de documentos basada en la NCD. La aplicación de las distancias basadas en compresión a la búsqueda de documentos no es trivial dado que este tipo de distancias tienen un punto débil que tiene que tenerse en cuenta si éstas se quieren aplicar en determinadas circunstancias. Su punto débil es que cuando los objetos comparados son muy diferentes en tamaño, las distancias obtenidas no son muy fiables. Un método de búsqueda de documentos que aborda este problema utilizando recuperación de pasajes se utiliza en la última parte de la tesis.

Los resultados experimentales muestran que la búsqueda de documentos se puede mejorar aplicando la técnica de distorsión presentada en la primera

parte de la tesis. Este hecho da mayor generalidad a los resultados obtenidos en la primera parte de la tesis, ya que dicha técnica ha resultado ser útil no sólo para el clustering de documentos, sino también para la búsqueda de los mismos [52].

Las principales contribuciones del Capítulo 7 se pueden resumir como sigue:

- Aplicación práctica de las principales conclusiones sacadas de los estudios llevados a cabo en las dos primeras partes de la tesis, a la búsqueda de documentos textuales.
- Mejora en la representación de los documentos que permite obtener un incremento considerable de la precisión en los resultados obtenidos al buscar dichos documentos.

Chapter 1

Introduction

Nowadays, most of the information stored electronically is stored in text form. In fact, if we think of the time that we spend every day reading e-mails, news, articles or reports, we will realize that most of the information that we use every day is text. This fact makes methods that deal with texts really interesting.

This thesis focuses on dealing with texts using compression distances. More specifically, it takes a step towards understanding both the nature of texts and the nature of compression distances.

The theoretical foundation of compression distances is the Kolmogorov complexity, which is intimately related to the concept of entropy proposed by Shannon in the paper that gave rise to Information Theory [113].

Broadly speaking, the theory developed by Shannon quantifies the amount of information as the amount of surprise that the information contains when revealed. A very simple way of understanding this is thinking of human communications.

For example, if one person tells another something that the latter already knows, there is no surprise in the message, and therefore, the first person has given the latter no information at all. On the contrary, if a person tells another something that the latter does not know, the first person has given the latter some information.

The amount of information transmitted in the second example, depends on the likelihood of the transmitted message. For example, saying “I have just looked out the window and I have seen a person walking in the street” gives less information than saying “I have just looked out the window and I have seen the Queen of England walking in the street”. Thus, the information defined by Shannon is inversely proportional to the probability, that is, the less probable a message is, the more information it contains.

Shannon defined the concept of entropy with the aim of formally quan-

tifying the information associated with a system. Entropy was defined as the average information gain from all possible events of the system. Given that each event can occur with a certain probability, the entropy created by Shannon weights the information associated with each event, according to the probability of the event.

The concept of entropy has been applied to numerous research areas. In particular, entropy is a basic concept in the area of data compression because it provides a theoretical bound on the amount of compression that can be achieved [7, 94, 103]. This theoretical bound coincides approximately with not only the entropy of the string, but also the Kolmogorov complexity of the string [123]. Therefore, the concept of entropy is directly related to the theoretical foundation of compression distances: the Kolmogorov complexity.

The Kolmogorov complexity of a string is defined as the length of the smallest program that can generate the string on a universal computer [66, 76]. A string would be more or less complex depending on its nature. For example, the string “0000000000000000” would be less complex than the string “0000111100001111”, and in turn, the latter would be less complex than the string “1011011100101010”.

The definition of Kolmogorov complexity can be extended to define the conditional Kolmogorov complexity, which measures the complexity of a string x relative to another string y . This measure is defined as the length of the smallest program that can generate the string x on a universal computer, having the string y as input to the program.

Li *et al.* defined a measure of similarity between two strings, called *Normalized Information Distance* -NID-, combining the concepts of Kolmogorov complexity, and conditional Kolmogorov complexity [75].

Given that Kolmogorov complexity is non-computable [123], NID is not computable either. However, Cilibrasi *et al.* proposed a computable measure, called *Normalized Compression Distance* -NCD-, that uses compression algorithms to estimate an upper bound upon the Kolmogorov complexity [30]. More detailed information on the NCD and the NID can be found in Section 4.3.

The NCD in particular and compression distances in general have been applied to several research areas because of their parameter-free nature, their wide applicability and their leading efficacy. Among others, they have been applied to document clustering [48, 49, 50, 51, 56, 121], document retrieval [52, 82], music classification [31, 46], data mining [32], security of computer systems [4, 12, 131], plagiarism detection [26], software engineering [3, 5, 109], bioinformatics [44, 65, 69, 91], chemistry [85], medicine [35, 107] or even art [119]. The fact that compression distances have been so widely used gives us an idea of how useful they are.

Despite their wide use, little has been done to interpret compression distances results or to explain their behavior. Whenever some analytical work on compression distances is carried out, it is usually focused on the algebraic manipulation of algorithmic information theory concepts [30, 75, 139].

One of the objectives of this thesis is to make progress on the understanding of compression distances in order to improve the performance of these metrics. In particular, this thesis focuses on one of the most important compression distances, the previously mentioned NCD. The analysis carried out in this thesis is mainly experimental. Therefore, the methodology used is the one used in experimental sciences. This methodology is based on disturbing the system to observe the consequences of the disturbance in the state of the system.

The assumption is that the information contained in the texts can be modified so that the compressor can better capture their structure, and therefore, the obtained NCD-based clustering results can be improved. The idea is to change the representation of the texts without losing relevant information so that this new representation is more suitable for compressors to better capture the similarities between the texts.

Before describing the experiments carried out throughout the thesis, Chapter 4 presents all the concepts needed to easily understand the contents of the thesis. After presenting them, Chapters 5 to 7 describe the experiments carried out throughout the thesis, and show the obtained experimental results. Each of these chapters has a clear objective, and generates a series of contributions, which are detailed always at the beginning of each chapter.

The research that corresponds to Chapter 5, tries to take a step towards the understanding of both the nature of textual information, and the nature of compression distances. This purpose is accomplished by analyzing how the information contained in the documents and how the upper bound estimation of their Kolmogorov complexity progress as words are removed from the documents. This is done by evaluating the impact that different distortion techniques, based on word removal, have on the NCD behavior [49, 50].

In particular, the research carried out in both Chapter 5 and Chapter 6 uses the NCD-based clustering method developed by the creators of the NCD [29], to measure the impact that the explored distortion techniques have on the NCD behavior. The use of the NCD-based clustering method as a tool to measure the performance of the NCD, allows analysis of how the information contained in the texts progresses as words are removed from the texts.

The main contributions of this chapter can be briefly summarized as follows:

- Analysis and study of new representations of texts to evaluate the be-

havior of the NCD.

- A technique to represent textual data, specially created to be used with compression distances, that reduces the complexity of the documents while preserving most of the relevant information.
- Experimental evidence of how to fine-tune the representation of texts to allow the compressor to obtain more reliable similarities and, therefore, to allow the compression-based clustering method to improve the non-distorted clustering results.

One of the main conclusions that can be drawn from the analysis made in Chapter 5, is that the accuracy of the clustering can be improved by applying a specific word removal technique. That technique implies, not only the removal of words, but also the maintenance of the previous text structure.

These results suggest that although the most important information of a text is contained in the most relevant words thereof, the information that surrounds these words is important too, because that information is the substrate that supports them. The hypothesis would be that this is the reason why the distortion technique that maintains the relevant information while preserving the contextual information is the best of all the evaluated distortion techniques.

Chapter 6 explores whether that hypothesis is correct or not. That is, the chapter studies how the results are affected by both the maintenance of the relevant information, and the maintenance of the contextual information.

The concept of contextual information has been used in several research areas. For example, it has been used in research areas such as contextual information retrieval [80, 100, 115, 117, 120], recommender systems [2, 71, 118, 132], context-aware computing applications [21, 42, 55, 96, 108], computer vision [1, 9, 27, 89, 90], speech recognition systems [43, 58, 74, 92] or network traffic analysis [47], among others.

In particular, in the management of textual data, the idea of context is very useful because it is strongly bound to texts due to their intrinsic nature. Since a text is not just a sequence of words, but it has coherent structure [67], applying the idea of context to text management arises naturally.

In this thesis, the contextual information is a byproduct of the application of the distortion technique presented in Chapter 5, that can improve the accuracy of the clustering. Chapter 6 compares that technique with three new distortion techniques created from it, which destroy the contextual information in different ways. Analyzing the obtained experimental results, it

can be observed that maintaining the contextual information is beneficial in NCD-based text clustering [48, 51].

The main contributions of Chapter 6 can be briefly summarized as follows:

- Experimental evaluation of the relevance that the contextual information has in compression-based text clustering, in a word removal scenario.
- New perspectives for the evaluation and explanation of the behavior of compression distances, in relation to contextual information.

Finally, Chapter 7, applies the knowledge acquired in Chapters 5 and 6 to NCD-based document search. The application of compression distances to document search is not trivial due to their having a weakness that must be taken into account if one wants to apply them under particular circumstances. Their drawback is that they do not commonly fit well when the compared objects have very different sizes. A document search method that addresses this issue by using passage retrieval is used in the last part of the thesis.

The experimental results show that the non-distorted document search results can be improved by applying the distortion technique presented in the first part of the thesis. This fact gives more generality to the results obtained in the first part of the thesis, since this technique has proven to be useful not only for document clustering, but also for document search [52].

The main contributions of Chapter 7 can be briefly summarized as follows:

- Practical application of the main conclusions taken from the studies developed in the first two parts of the thesis to document search.
- Improvement in the representation of documents that allows increasing the accuracy of the results obtained when searching documents.

Chapter 2

Objectives

Broadly speaking, this thesis applies text distortion to compression-based text clustering with the aim of taking a step towards understanding the nature of compression distances, and the nature of textual data. After that, it applies text distortion to compression-based document retrieval with the aim of exploring a possible practical application of the knowledge acquired in the first study. These widespread objectives can be divided into more specific goals:

- *Objective 1. Providing new perspectives for understanding the nature of textual data.*

The huge amount of information stored in text form makes the study of the nature of texts really interesting. Many research areas address several aspects of processing textual information in different manners. This thesis uses compression distances to explore how the application of different distortion techniques affects the information contained in the evaluated texts.

- *Objective 2. Providing a technique to smoothly reduce the complexity of the documents while preserving most of their relevant information.*

Removing irrelevant parts of the data has been found to be beneficial in data analysis. In fact, most of the research areas that work with textual data apply that idea to text processing. This thesis tries to provide a text distortion technique that reduces the complexity of the texts while maintaining most of the relevant information contained in them.

- *Objective 3. Giving experimental evidence of how to fine-tune the text representation so that better results are obtained when using NCD-driven text clustering.*

One of the purposes of this thesis is finding a text representation that can improve the clustering results. This work explores different distortion techniques with the aim of attaining this objective. One of the explored techniques has been found to be beneficial to NCD-based text clustering.

- *Objective 4. Giving new insights for the evaluation and explanation of the behavior of the NCD.*

Compression distances have been widely used in knowledge discovery and data mining due to their parameter-free nature, wide applicability, and leading efficacy in several domains. However, little has been done to interpret their results or to explain their behavior. This thesis tries to shed light on this issue by performing an experimental study on text distortion.

- *Objective 5. Experimentally evaluating the relevance that the contextual information has in compression-based text clustering, in a word removal scenario.*

The distortion technique that fine-tunes the text representation implies not only the removal of words, but also the maintenance of the previous text structure. Exploring the relevance of both factors becomes necessary in order to better understand the results. This research is carried out in the thesis as well.

- *Objective 6. Applying the main conclusions taken from the studies developed in the first two parts of the thesis to document search.*

A problem in which the distortion technique that fine-tunes the text representation can be very useful is the search of texts. Applying that technique to document search is one of the purposes of this work.

- *Objective 7. Giving a representation of documents that improves the non-distorted document search accuracy.*

Text representation plays an important role in document search. Thus, good text representations can improve the accuracy of the results, whereas bad ones can make the results get worse. Exploring if the application of the above mentioned distortion technique can lead to better document search results is the final goal of this thesis.

Chapter 3

Thesis Overview

The thesis is structured as follows:

- *Chapter 4* presents and discusses all the concepts needed to easily understand the contents of the thesis.
- *Chapter 5* explores several text distortion techniques based on word removal. It analyzes how the information contained in the documents and how the upper bound estimation of their Kolmogorov complexity progress as the words are removed from the documents in different manners.
- *Chapter 6* explores how the loss or the maintenance of the contextual information affects the clustering accuracy. At the same time, it explores how the loss or the preservation of the remaining words structure affects the clustering.
- *Chapter 7* applies the distortion technique that can lead to better clustering results to document search.
- *Chapter 8* discusses the conclusions drawn from the research carried out in the thesis.
- *Chapter 9* summarizes each of the contributions made from the work developed throughout the thesis.
- *Chapter 10* presents the papers created from the investigation carried out in the thesis.
- *Appendix A* presents an index of the acronyms used.

- *Appendix B* contains the detailed description of the datasets used throughout the thesis. In addition, it shows, as a sample, a fragment of a document for each dataset.
- *Appendix C* contains the detailed description of the queries used in the experiments carried out in Chapter 7. It also shows, as a sample, a fragment of a query for each dataset.
- *Appendix D* contains all the detailed results obtained in the work presented in Chapter 5.

Chapter 4

Related Work

This chapter presents and discusses all the concepts needed to easily understand the contents of the thesis.

Compression distances have to be described in this chapter because this thesis uses them to cluster and retrieve documents. Since, compression distances are based on information theory concepts, the latter have to be presented as well, in order to help and understand compression distances. Furthermore, given that compression distances use compression algorithms to calculate the similarity between two objects, the compression algorithms explored in this thesis have to be described.

Three compression algorithms are used in this thesis to calculate compression distances. Each of them belongs to a different family of compressors: LZMA, PPMZ, and BZIP2. LZMA compressor, is a Lempel-Ziv-Markov chain algorithm [97]. PPMZ compressor is an adaptive statistical data compression algorithm based on context modeling and prediction [13]. BZIP2 compressor is a block-sorting compressor based on the Burrows-Wheeler Transform, Huffman codes, the Move-To-Front transform, and Run Length Encoding [19, 59, 103, 111]. Among all the existing compression algorithms, only these ones are reviewed in this chapter.

In addition, given that the distortion techniques explored in this thesis are based on word removal, the concept of word removal must be presented as well. Moreover, since the most important distortion technique used in the thesis maintains the contextual information despite the removal, presenting how several research areas apply the concept of contextual information is necessary.

The chapter contains a section for each of the concepts mentioned above.

4.1 Information Theory Concepts

Information Theory -IT- is a branch of applied mathematics and electrical engineering that focuses on the task of quantifying information. The famous work by Claude Shannon in 1948 [113] involved its creation. The research area of IT has turned out to be one of the most influential ones because of its wide applicability in many other domains [123].

Roughly speaking, the theory developed by Shannon quantifies the amount of information as the amount of surprise that the information contains when revealed. A very simple way of understanding this is thinking of human communications. For example, if one person tells another something that the latter already knows, there is no surprise in the message, and therefore, the first person has given the latter no information at all. On the contrary, if a person tells another something that the latter does not know, the first person has given the latter some information.

The amount of information transmitted in the second example, depends on the likelihood of the transmitted message. For example, saying “I have just looked out the window and I have seen a person walking in the street” gives less information than saying “I have just looked out the window and I have seen the Queen of England walking in the street”. Thus, the information should be proportional to the probability, that is, the less probable a message is, the more information it contains. The mathematical formulation of this idea would be:

$$I(x) \sim \frac{1}{P(x)} \quad (4.1)$$

where x is the event, and $P(x)$ is the probability function.

Furthermore, independent information should be additive, that is, if the first person tells the second one something more, then the first has given the latter some more information, independent of, and additional to, the information that the former gave the latter previously.

Since the probability of independent events is the product of the probabilities of the individual events, the function used to represent the information should have the following property:

$$f(xy) = f(x) + f(y) \quad (4.2)$$

The mathematical function that transforms a product in an addition is the logarithm. That is the reason why logarithms are used to give a measure of the information.

Therefore, from all the above, the following formal definition can be derived:

$$I(x) = \log_2 \frac{1}{P(x)} \quad (4.3)$$

Several examples will be analyzed in order to better understand the basis of IT. The simplest event that can be analyzed in order to approach the basis of IT is the toss of a coin. Before the toss, the result is uncertain, this uncertainty must be resolved by tossing the coin. This will produce either a head or a tail. Thus, the result of tossing the coin can be expressed with a single bit since there are only two possibilities. Therefore, the information contained in the result is one bit.

This strategy can easily be generalized to resolve more complex problems. The idea is finding the minimum number of yes/no questions that must be answered in order to resolve the uncertainty. The number of questions will correspond to the number of bits needed to express the information in the result because the information contained in a yes/no question is a bit, since this kind of question only produces two possible answers.

A more complex problem that intuitively introduces why the logarithm is the mathematical function that quantifies information, is drawing a card from a deck of 32 playing cards. This event can be thought as guessing a number between 1 and 32. The minimum number of yes/no questions needed to guess a number between 1 and 32 is given by the binary search algorithm.

In computer science, a binary search locates an element in a sorted array of elements. The algorithm works by comparing the searched element with the element contained in the middle of the array. The comparison determines whether the element is already found or must be searched for again in the left half of the array or in the right half of it. The asymptotical cost of this algorithm is $\log_2 N$, N being the number of elements contained in the array. This reasoning constitutes an alternative way of explaining why the logarithm is the mathematical function that quantifies information.

Returning to the problem of guessing the card, the number of questions that have to be answered to guess the card is 5, because $\log_2 32 = 5$.

One can easily interpret that result thinking of the bits needed to codify the number and the suit of the card from a deck of 32 playing cards. Since there are four possible suits, two bits will be required to codify the suit. Similarly, since there are eight possible numbers, three bits will be necessary to codify the number. Again, this makes a total of 5 bits to codify the card:

- *Suit of the card:* 2 bits because there are 4 possible suits.
- *Number of the card:* 3 bits because there are 8 possible numbers.

To quantify the information associated with a system, Shannon defined the concept of entropy as the average information gain of all possible system events [113]. Since each event can occur or not, with a certain probability, the entropy gives a weight to the information associated with each event, according to the probability of the event. Mathematically, the entropy $H(X)$ of a discrete random variable X , with probability function $p(x)$ is defined as follows:

$$H(X) = - \sum_{x \in \chi} p(x) \log_2 p(x) \quad (4.4)$$

Note that the entropy of X can also be interpreted as the expected value of $\log_2 \frac{1}{p(X)}$.

The expected value of a random variable $g(X)$ is as follows:

$$E_p g(X) = \sum_{x \in \chi} p(x) g(x) \quad (4.5)$$

Therefore:

$$E_p \log_2 \frac{1}{p(X)} = \sum_{x \in \chi} p(x) \log_2 \frac{1}{p(x)} = \sum_{x \in \chi} p(x) \log_2 p(x)^{-1} = - \sum_{x \in \chi} p(x) \log_2 p(x)$$

Thus:

$$H(X) = E_p \log_2 \frac{1}{p(X)} \quad (4.6)$$

The simple example discussed above, the toss of a coin, can be used to clarify the concept of entropy. The probability of obtaining a head or a tail is the same:

$$X = \begin{cases} head & \text{with probability } \frac{1}{2} \\ tail & \text{with probability } \frac{1}{2} \end{cases}$$

Then,

$$H(X) = - \sum_{x \in \chi} p(x) \log_2 p(x)$$

$$H(X) = - \left[\frac{1}{2} \log_2 \frac{1}{2} + \frac{1}{2} \log_2 \frac{1}{2} \right] = - \log_2 \frac{1}{2} = - \log_2 2^{-1} = \log_2 2 = 1 \text{ bit.}$$

Similarly, the entropy of the example of drawing a card from a deck of 32 cards can be calculated, given that the probability of drawing a card is the

same for all the cards. That is: $p(x) = \frac{1}{32}$.

$$H(X) = - \sum_{x \in \mathcal{X}} p(x) \log_2 p(x) = - [32 (\frac{1}{32} \log_2 \frac{1}{32})] = \log_2 32 = 5 \text{ bits.}$$

Notice that these examples have an important characteristic: all the possible events have the same probability of occurrence. In general, in systems in which all the possible events have the same probability of occurrence, the entropy is equivalent to the logarithm of the number of possible events. Thus if N is the number of possible events:

$$H(X) = - \sum_{i=1}^N \frac{1}{N} \log_2 \frac{1}{N} = \log_2 N \text{ for equally likely events.} \quad (4.7)$$

Let us analyze a more complex example. The setting is the same as the previous example, that is, drawing a card from a deck of 32 cards. However, in this example, the amount of uncertainty of an event E given another event F is calculated. For example, given the following events:

- E = The card drawn is the ace of hearts.
- F = The card drawn is a heart.

The probability of E given F is:

$$P(E/F) = \frac{P(E \cap F)}{P(F)} = \frac{P(E)}{P(F)} \text{ as } E \subset F$$

The probabilities of the events E and F are:

- $P(E) = \frac{1}{32}$, since there is only one ace of hearts in the deck.
- $P(F) = \frac{1}{4}$, since there are four suits in the deck.

Therefore:

$$H(E/F) = - \log_2 P(E/F) = \log_2 \frac{P(E)}{P(F)} = \log_2 \frac{32}{4} = \log_2 8 = 3 \text{ bits.}$$

This result can be easily interpreted. The fact that F has occurred determines the suit of the card, that is, determines two bits, because as said previously, two bits are needed to codify the suit of the card because there are four possible suits. Consequently, specifying the card given that it is a heart, requires only $5 - 2 = 3$ bits. Thus, the uncertainty of E has been reduced thanks to the knowledge of F .

The main theorem proved by Shannon says that a message of n symbols can, on average, be compressed down to nH bits, but not further. It also says that almost optimal compressors -called *entropy encoders*- exist [103].

Directly related to the measure of information proposed by Shannon is the Kolmogorov complexity of a string x , $K(x)$. Andréi Kolmogórov defined the algorithmic complexity of an object x , $K(x)$ as the length of the shortest program that can generate x on a universal computer [66, 76]. This definition can be extended to define the conditional Kolmogorov complexity. That is, the Kolmogorov complexity of a string x , relative to another string y . The conditional Kolmogorov complexity $K(x|y)$ is the length of the smallest program that generates the string x having the string y as input to the program.

The best way to assimilate the concept of Kolmogorov complexity is intuitively analyzing some strings:

- The question is: what is the shortest program that can generate each of these strings?

The second string can be described as a “11”, followed by r_i repetitions of “0”, where r_i can be 2, 4 or 6. Therefore, the shortest program that can generate such a string is more complex than the previous one.

The good news is that most binary strings used in practice to represent texts are similar to the second string shown previously. Therefore, they exhibit some regularity, and thus they can be compressed [103].

The concept of Kolmogorov complexity is directly related to this thesis because it has been used to define a measure of similarity between two strings, giving rise to the concept of *Normalized Information Distance* -NID- [75]. The compression distance used in this thesis is created from the NID. Both distances are described in depth in Section 4.3.

4.2 Compression Algorithms

Data compression existed before the appearance of computers, as some well-known codes, such as the Braille of 1825 or the Morse of 1838 show. Interesting approaches were used in both cases, as explained below.

The Braille code is based on a communication method developed by Charles Barbier in order to allow Napoleon's soldiers to communicate silently and lightlessly. The Barbier method was rejected by the military due to that it encoded each letter with a set of 12 embossed dots, making it too difficult for soldiers to read by touch. However, it ended up leading to the creation of the extremely important Braille code that has allowed blind people to read since its creation.

The inception of the Braille code is due to the encounter between Charles Barbier and Louis Braille in the National Institute for the Blind in Paris in 1821. Braille, who only was 12 years old, associated the failure of the method to the high number of dots used to encode each letter. His hypothesis was that since the human fingertip could not cover the whole symbol without moving, the message could not be read efficaciously and efficiently. This led to the creation of the Braille system, which encodes each symbol with 6 dots.

Each of the 6 dots in a symbol can be flat or raised, which means that the information contained in a symbol is equivalent to 6 bits, which implies the possibility of coding $2^6 = 64$ different symbols. Since the letters, digits, and punctuation marks do not require the use of all the codes, the spare ones are used to code common words, such as *and*, *for* and *of*, and common strings of letters, such as *ound*, *ation*, and *th*. Although this kind of data compression is modest, it is important because books in Braille are usually very large due to the room that each symbol takes up.

The first version of the Morse code, mentioned above, which dates from 1832, allows the transmission of textual information as a series of short and long dashes that represent numbers. A code book or dictionary associates each number with a word. Thus, this first version of the Morse code was a primitive form of data compression.

The famous Morse code used nowadays is the evolution of the primitive Morse code. It allows the transmission of textual information as a series of dots and dashes as well. It encodes letters, digits, and some punctuation marks, using variable-size codes to encode each symbol. This important feature of the Morse code leads to a better efficiency because the length of each symbol is approximately inversely proportional to its frequency of occurrence in English. This is reminiscent of the basic idea of the Huffman coding, which will be considered later.

These are some examples of data compression used before the appearance

of computers. After that, in the computer age, data compression has become crucial, initially, to reduce the storage needed for data, and later, after the appearance of the Internet, to reduce transmission time.

Many compression strategies have been used since the emergence of data compression as a research field, from primitive algorithms to sophisticated algorithms that achieve very high compression rates. Among these latter, most text compression methods are either dictionary or statistical based. The next subsections explain in more detail the characteristics of the most important text compression algorithms.

4.2.1 Statistical Methods

Statistical compressors are based on developing statistical models of the text. The model assigns probabilities to the input symbols, and then, the symbols are coded based on these probabilities. The model can be *static* or *dynamic* -also known as *adaptive*-.

Huffman Coding

David Huffman developed this entropy encoding algorithm in 1952 [59]. This method uses variable-length codes for encoding the symbols using bits. It assigns shorter codes to the more frequent symbols and longer codes to the less frequent ones to make the coding more efficient.

The method constructs a binary tree, with a symbol at each leaf, which can be traversed to determine the codes of the symbols. Fig 4.1 shows an example of a Huffman tree generation.

The process is as follows:

1. A list of nodes that contains the alphabet symbols is created and sorted in increasing order of frequency.
2. Then, the tree is constructed from that list following these steps:
 - (a) Remove the two nodes of lowest frequency from the list.
 - (b) Create a new node with these nodes as children and with frequency equal to the sum of the children's frequencies.
 - (c) Insert the new node into the ordered list of nodes.
3. At the end of the process, a binary tree, which has a leaf for each symbol of the alphabet, is obtained.

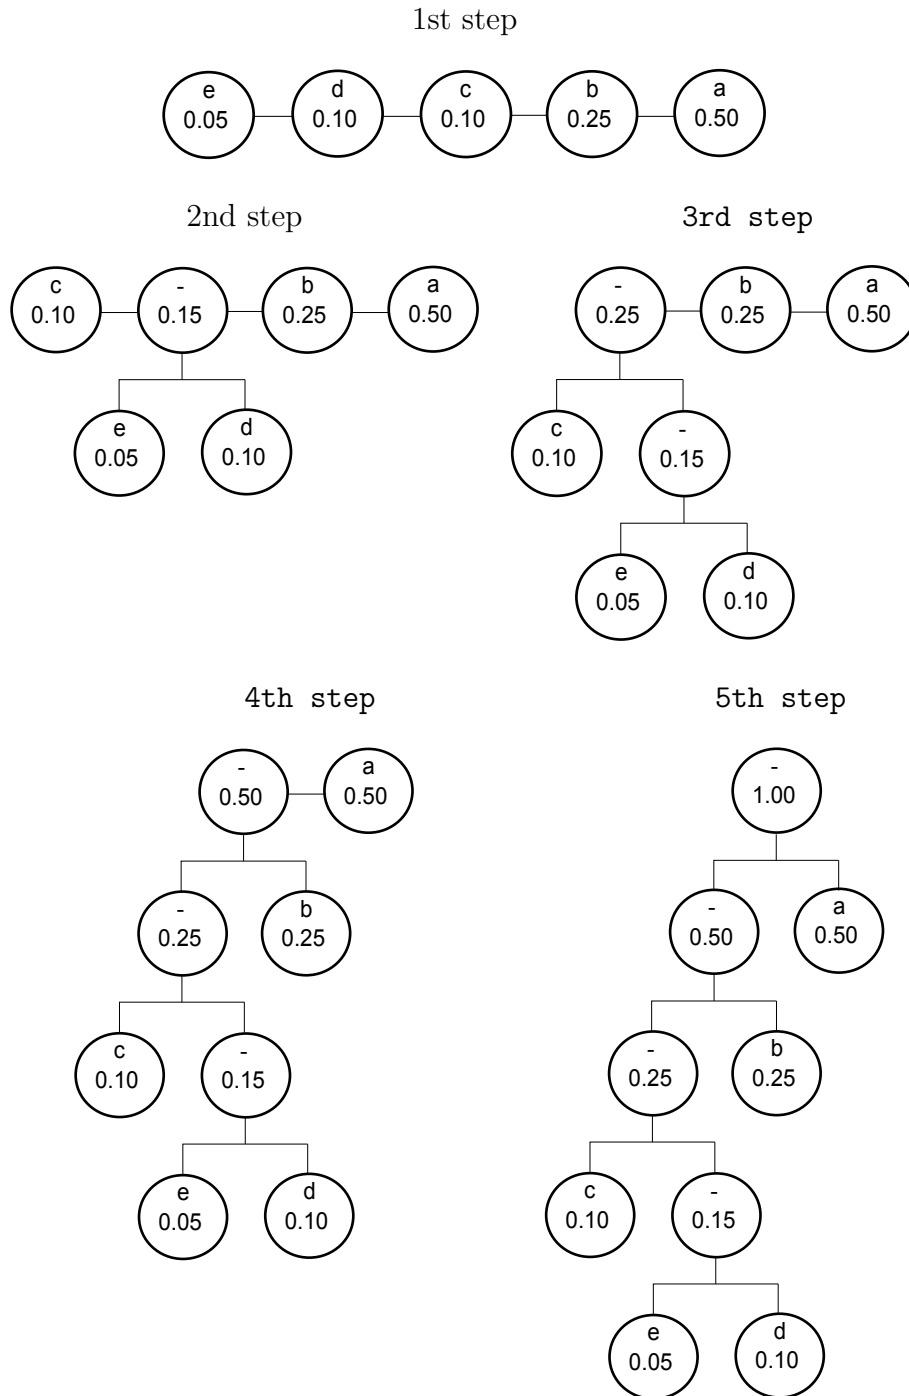


Figure 4.1: Huffman tree generation. The tree is constructed from the list of nodes shown in the 1st step. That list contains the alphabet symbols and their frequencies sorted in increasing order of frequency. In each step, the list is updated by removing its first two nodes, and then by inserting a new node that has these nodes as children.

This binary tree is then used to assign the codes to the symbols by traversing the tree from the root node to the leaf that contains the symbol that is being coded. Since the tree is binary, there are two possibilities of going from one node to the next one in the tree traversal process. One is going through the left child of the node, while the other is going through the right one. The coding process assigns a different bit in every step depending on the edge used to go from one level to the next. This implies that the Huffman coding results in a prefix code, due to the fact that the bit string representing some particular symbol is never a prefix of the bit string representing any other symbol.

PPM

The PPM algorithm, whose name stands for *Prediction with Partial string Matching*, is an adaptive statistical data compression technique based on an encoder that maintains a statistical model of the text. It was originally developed by John Cleary and Ian Witten [34], with extensions and an implementation by Alistair Moffat [87].

There can be many statistical models depending on the way the input data is treated. Thus, statistical models can take into account separated symbols or groups of contiguous symbols. While the former do not consider the context of the symbols because they treat them separately, the latter do consider it because they take into account the preceding symbols of each symbol. Because of that, they receive the name of *context-based* statistical models.

Depending on whether the probabilities are fixed or dynamic, that is, updated as more data is being input, the modeler would be *static* or *dynamic* -also known as *adaptive*-. The latter are more suitable because they adapt to the particularities of the data contained in the file being compressed.

Although in principle, it can seem logical that a long context is better than a small one because the longer retains information about the nature of old data, experience shows that large data files contain different distributions of symbols in different parts. Thus, better compression can be achieved if the model takes into account contexts of about 10 symbols [103].

In general, an order- N adaptive context-based modeler considers the N symbols preceding the symbol being processed. Although this approach may sound good, there is a problem with it. The drawback is that considering only order- N contexts can lead to no compression in spite of the existence of smaller order instances which could be used to compress the data. That is, when the encoder does not find any order- N instance of a given symbol, it simply writes the symbol on the compressed stream as a literal. However, the

data could be compressed using smaller contexts. The PPM method solves this problem by switching to shorter contexts if necessary. Thus, the PPM method uses smaller and smaller parts of the context in order to achieve a better compression.

PPM uses sophisticated data structures and it usually achieves the best performance of any real compressor although it is also usually the slowest and most memory intensive [30]. One of the data structures that can be used to implement the PPM algorithm is a special type of tree called *trie*.

Level 1 of a *trie* contains the order-1 contexts, which means that it contains one node for each symbol read so far. Level 2 contains all the order-2 contexts, and so on. In a *trie*, each context can be found by traversing the tree from the root to one of the leaves.

Fig 4.2 illustrates an example that helps to understand the process of creation of a *trie* and the meaning of the nodes contained in it. The figure shows the seven steps needed to construct the *trie* for the string “bananas”, assuming $N = 2$. Note that the tree grows in width but not in depth. In fact, it can be observed that its depth remains $N + 1$ regardless of how many characters have been read.

The characters in the string are processed first to last, one at a time. All the intermediate *tries* shown in the figure illustrate the state of the *trie* after processing each character. The numbers in the nodes are context counts. Notice that three nodes are involved in each step, except the first two steps when the *trie* has not yet reached its final height. All the nodes involved in each step are shaded to ease the understanding of the figure.

The first tree contains only one node because only one character (“b”) has been processed so far. The label “b,1” on the node means that the “b” has occurred only once until that moment.

After reading the next symbol of the string (“a”), the tree is updated by adding two nodes. The “a,1” on level 1 means that the character “a” has occurred only once. The “a,1” that is on level 2, under the “b,1”, means that the substring “ba” has occurred only once.

After reading the next symbol (“n”), the tree is updated by adding three nodes, one at each context:

- The node “n,1” on level 1 means that the character “n” has occurred once.
- The node “n,1” on level 2 means that the substring “an” has occurred once.
- Finally, the node “n,1” on level 3 means that the substring “ban” has been seen once.




Figure 4.2: PPM: seven *tries* of “bananas” for a context of $N = 2$. The characters are processed first to last. The numbers in the nodes are context counts. Three nodes are involved in each step, except the first two steps when the *trie* has not yet reached its final height. The nodes involved are shaded to ease the understanding of the figure. For example, after reading the first “n”, the tree is updated by adding three nodes. The node “n,1”, on level 1, means that the character “n” has occurred once, so far. The node “n,1”, on level 2, means that the sequence “an” has occurred once. Finally, the node “n,1”, on level 3, means that the sequence “ban” has occurred once.

The last *trie* of Fig 4.2, that is, the 7th one, can be analyzed to see the contexts that correspond to the string “bananas”. For example:

- The “a,3” on level 1 of the tree means that the “a” occurs 3 times in the string “bananas”:
 - bananas
 - bananas
 - bananas
- The “n,2” and “s,1” below it mean that these three occurrences of “a” were followed by “n” twice, and by “s” once:
 - bananas
 - bananas
 - bananas
- These two occurrences of “an”, were followed always by “a”, as the node “a,2” on level 3 indicates.
 - bananas
 - bananas

Many variants of the PPM algorithm have been implemented [103]: PPMA, PPMB, PPMP, PPMX, PPMZ. However, the bases of the method are always the ones explained above.

In this thesis, the variant called PPMZ is used. The PPMZ algorithm, implemented by Charles Bloom [15], tries to improve the PPM performance by handling features such as deterministic contexts, unbounded-length contexts, and local order estimation, in an optimal way [103]. Implementation details are difficult to understand due to the code being very obscure. However, since PPMZ belongs to the family of PPM algorithms, the basis of it are the ones explained above.

4.2.2 Dictionary Methods

Dictionary compressors break the text into fragments that are saved in a data structure called *dictionary*. When a fragment of new text is found to be identical to one of the dictionary entries, a pointer to that entry is written on the compressed stream.

The simplest example of a dictionary compressor can be one that uses an English dictionary to compress English texts, by coding each word as its index in the dictionary, or by writing the word into the output stream when the word is not found in the dictionary. Obviously, this kind of approach is not a good choice for a general-purpose compressor since the words contained in the dictionary do not depend on the input.

The most famous dictionary compressors are the ones that belong to the Lempel-Ziv family [103]. The origin of this family of compressors is the LZ77, also known as LZ1, and the LZ78, also known as LZ2, which were developed by Jacob Ziv and Abraham Lempel [143, 144].

LZ77

This algorithm uses as dictionary part of the input stream previously seen. The method is based on a sliding window that the encoder shifts as the strings of symbols are being encoded. That is the reason why sometimes this method is called *sliding window*.

The window is divided into two parts, the first part, called the *search buffer*, is the current dictionary, while the second part, called the *look-ahead buffer* contains the text yet to be encoded. It is important to point out that practical implementations of this method use really long *search buffers* of thousands of bytes long, and small *look-ahead buffers* of tens of bytes long [103].

The encoding algorithm works as follows:

1. It scans the *search buffer* backwards looking for a match to the first symbol in the *look-ahead buffer*.
2. Then, it calculates the length of the match by comparing the symbols following the symbol found.
3. After that, it keeps doing this in order to find longer matches.
4. After the search process, it selects the longest one, or the last one found in the event of a tie. This is done this way to avoid having to memorize previously found matches.
5. Finally a token with three parts -offset, length of match, and first symbol in the *look-ahead buffer*- is written on the output in this way:
 - (a) If the backward search yields no match, a token with zero offset, zero length, and the unmatched symbol is written on the output. Then, the window is shifted to the right one position.

- (b) If there is a match, a token with the offset, the length of match, and the symbol that follows the matched sequence in the *look-ahead buffer* is written on the output. Then, the window is shifted to the right $L + 1$ positions, L being the length of match.

To sum up, the LZ77 encodes the input by generating tokens with three parts: offset, length and next symbol in the *look-ahead buffer*. Table 4.1 shows an example that helps to understand the algorithm. It shows the evolution of the *search buffer* and the *look-ahead buffer* for the input data “the-abbess-and-the-abbot-are-in-the-abbey”.

Let us analyze some steps of the process to ease the understanding of the encoding algorithm. Since the *search buffer* is empty at the beginning of the process, the first token is (0,0,'t') because the backward search yields no match, and the unmatched symbol is the character 't'. In fact, the first six tokens have an offset and a length of 0 because the first six characters of the input data are different.

After processing the first six characters, there is a match of offset 1 and length 1 because the last character in the *search buffer* and the first character in the *look-ahead buffer* are the same ('b'):

- search buffer: “the-abb”
- look-ahead buffer: “bess-and-the-abbot-are-in-the-abbey”

This explains why the seventh token written on the output is (1,1,'e'). Note that including the character 'b' in the token is not necessary because it is implicitly included thanks to the offset and the length of 1.

A more interesting circumstance occurs after processing the first 16 characters of the input. It is easy to see that at that point, there is a match of length 6 at a distance of 15, as can be observed looking at the content of the buffers:

- search buffer: “the-abbess-and-t”
- look-ahead buffer: “he-abbot-are-in-the-abbey”

This explains why the thirteenth token written on the output is (15,6,'o').

In this thesis, the Lempel-Ziv-Markov chain algorithm LZMA [97], created by Igor Pavlov, is used. This is a compression algorithm that uses a variant of the LZ77 to encode the input, and then uses a range encoder to encode the output obtained by the LZ77.

Table 4.1: LZ77: Lempel-Ziv sliding window. The algorithm scans the *search buffer* backwards looking for a match to the first symbol in the *look-ahead buffer*. It keeps doing this in order to find the longest match. Then, it selects the longest one, or the last one found in the event of a tie. Finally, a token with the offset, the length of match, and the first symbol in the *look-ahead buffer* is written on the output.

Search buffer	Look-ahead buffer	Token
	the-abbess-and-the-abbot-are-in-the-abbey	(0,0,'t')
t	he-abbess-and-the-abbot-are-in-the-abbey	(0,0,'h')
th	e-abbess-and-the-abbot-are-in-the-abbey	(0,0,'e')
the	-abbess-and-the-abbot-are-in-the-abbey	(0,0,'-')
the-	abbess-and-the-abbot-are-in-the-abbey	(0,0,'a')
the-a	bess-and-the-abbot-are-in-the-abbey	(0,0,'b')
the-ab	ess-and-the-abbot-are-in-the-abbey	(1,1,'e')
the-abbe	ss-and-the-abbot-are-in-the-abbey	(0,0,'s')
the-abbes	s-and-the-abbot-are-in-the-abbey	(1,1,'-')
the- <u>abbess</u> -	and-the-abbot-are-in-the-abbey	(7,1,'n')
the-abbess-an	d-the-abbot-are-in-the-abbey	(0,0,'d')
the-abbess-and	-the-abbot-are-in-the-abbey	(11,1,'t')
the- <u>abbess</u> -and-t	he-abbot-are-in-the-abbey	(15,6,'o')
the-abbess-and-the-abbo	t-are-in-the-abbey	(23,1,'-')
the- <u>abbess</u> -and-the-abbot-	are-in-the-abbey	(21,1,'r')
the-abbess-and-the-abbot-ar	e-in-the-abbey	(25,2,'i')
the-abbess-and-the-abbot-are-i	n-the-abbey	(18,1,'-')
the- <u>abbess</u> -and-the-abbot-are-in-	the-abbey	(32,8,'y')
the-abbess-and-the-abbot-are-in-the-abbey		

Range encoding is a data compression technique created by G. Nigel N. Martin [81] that encodes all the symbols of the message into one number using a probability estimation.

The LZMA produces a stream of literal symbols and phrase references, which is encoded one bit at a time by the range encoder, using a model to make a probability prediction of each bit. This gives much better compression because it avoids mixing unrelated bits together in the same context. In fact, empirical evidence shows that it performs very well on structured data and it looks very much like any other LZ algorithm. However, it trounces them all [14].

4.2.3 Other Methods

Some modern context-based text compression methods perform a transformation on the input data and then apply a statistical model to assign probabilities to the transformed symbols.

BZIP2

BZIP2 is a block-sorting compressor developed by Julian Seward [111]. BZIP2 compresses data using Run Length Encoding, the Burrows-Wheeler Transform, the Move-To-Front transform and Huffman coding.

The algorithm reads the input stream block by block and each block is compressed separately as one string. The length of the blocks is between 100 and 900 KB. The compressor uses the Burrows-Wheeler Transform to convert frequently-recurring character sequences into strings of identical letters, and then it applies Move-To-Front transform and Huffman coding. All these methods are explained later so the basis of the BZIP2 compressor can be understood.

Run Length Encoding

Run Length Encoding -RLE- is a very simple form of data compression in which, if a data item d occurs n consecutive times in the input stream, the occurrences are replaced with the single pair nd . The sequences in which the same data value occurs in many consecutive data elements are called a *run length* of n .

The main problem with this method is that, in plain English texts, there are many sequences of two equal symbols but a sequence of three is rare. However, this method can be combined with other methods to process the text before RLE so the new text representation is more suitable to achieve bigger compression rates. This is precisely what the BZIP2 compression algorithm does, because it applies RLE after applying the Move-To-Front transform.

Move-to-Front

The Move-To-Front transform -MTF- [11, 102] is an encoding of data usually used as an extra step in data compression algorithms, such as for example BZIP2. Table 4.2 shows an example that helps to understand how the MTF transform works.

The method transforms the data into a sequence of integers in the following manner. It maintains a list that stores the symbols of the alphabet in such a way that the most frequent ones are maintained near the front. This is done by updating the list each time a symbol is processed, moving it to the front. Then, a symbol is encoded as the number of symbols that precede it in the list, or in other words, it is encoded as its index in the list, 0 being the index of the first element.

This implies that long sequences of identical symbols are replaced by many zeros, and frequently used symbols are coded with small numbers. The MTF transform takes advantage of local correlation of frequencies to reduce the entropy of a message. In other words, when the characters exhibit local correlations, the sequence of integers will contain small numbers [103].

The MTF transform is used in the Burrows-Wheeler Transform, because the latter is very good at producing a sequence that exhibits local frequency correlation from text.

Table 4.2: Move-To-Front transform. The algorithm transforms the data into a sequence of integers. It maintains a list of the symbols of the alphabet in such a way that the most frequent ones are maintained near the front. In order to do so, the list is updated every time a symbol is processed, moving it to the front. A symbol is encoded as the number of symbols that precede it in the list.

Iteration	Output	List
pebblepebble		(abcdefghijklmnopqrstuvwxyz)
<u>p</u> ebblepebble	15	(pabcdefghijklmnopqrstuvwxyz)
pe <u>b</u> blepebble	15,5	(epabcdefghijklmnopqrstuvwxyz)
peb <u>b</u> lepebble	15,5,3	(bepabcdefghijklmnopqrstuvwxyz)
pebble <u>p</u> ebble	15,5,3,0	(bepabcdefghijklmnopqrstuvwxyz)
pebblep <u>e</u> bble	15,5,3,0,12	(lbepabcdefghijklmnopqrstuvwxyz)
pebblepe <u>b</u> ble	15,5,3,0,12,2	(elbepabcdefghijklmnopqrstuvwxyz)
pebblepeb <u>b</u> le	15,5,3,0,12,2,3	(pelbepabcdefghijklmnopqrstuvwxyz)
pebblepebb <u>e</u>	15,5,3,0,12,2,3,1	(eplbepabcdefghijklmnopqrstuvwxyz)
pebblepebb <u>p</u> le	15,5,3,0,12,2,3,1,3	(beplabcdefghijklmnopqrstuvwxyz)
pebblepebbp <u>e</u> le	15,5,3,0,12,2,3,1,3,0	(beplabcdefghijklmnopqrstuvwxyz)
pebblepebbp <u>e</u> bble	15,5,3,0,12,2,3,1,3,0,3	(lbepabcdefghijklmnopqrstuvwxyz)
pebblepebbp <u>e</u> b <u>b</u> le	15,5,3,0,12,2,3,1,3,0,3,2	(elbepabcdefghijklmnopqrstuvwxyz)

Burrows-Wheeler Transform

The Burrows-Wheeler Transform -BWT- is an algorithm created by Michael Burrows and David Wheeler [19] that is applied by the BZIP2 compressor.

BWT permutes the order of the characters of the string being transformed with the purpose of bringing repetitions of the characters closer. This is useful for compression, since there are techniques such as MTF and RLE that work very well when the input string contains runs of repeated characters.

Although in practice the BWT implementation is more complex than the

algorithm explained below, this version can be more easily understood while keeping the same philosophy as the complex one.

Table 4.3 shows how the algorithm works when it is used to encode the string “sentence”.

The algorithm works as follows:

1. The encoder creates an $n \times n$ matrix. It stores the string to code in the first row. The rest of the rows contain $n - 1$ copies of the said string, each cyclically shifted one symbol to the left.
2. Then the matrix is sorted lexicographically by rows.
 - Notice that the last character of a row is always the one that precedes the first character in that row.
 - Notice too, that every row and every column of the matrix is a permutation of the string being transformed.
3. Finally, the last column of the sorted matrix is taken as the transformed version of the input string.

Applying this algorithm creates more easily compressible data, because sorting the rotations of the string tends to create regions that concentrate just a few symbols. However, the BWT works well only if the length of the string is large -at least several thousand symbols per string- [103].

The only information needed to reconstruct the original string from the last column, is the row number of the original string in the lexicographically sorted matrix. Thus, the decoding process works thanks to these facts:

1. The encoded string, contains all the characters in the text. Therefore, it can be used to get the first column of the lexicographically sorted matrix by simply sorting the encoded string.
2. Since the last character of a row is always the one that precedes the first character in that row, and given that the first and the last column of the matrix are held, both columns can be used to obtain all pairs of successive characters in the original string, where pairs are taken cyclically so that the last and first character form a pair.
3. After reconstructing the lexicographically sorted matrix, the original string can be obtained from the row number of the original string in the sorted matrix.

Table 4.3: Burrows-Wheeler Transform encoding. The algorithm stores the string to code in the first row of the cyclically shifted matrix. The rest of the rows contain $n - 1$ copies of the said string, each cyclically shifted one symbol to the left. Then, the lexicographically sorted matrix is created by sorting the said matrix by rows. Finally, the last column of the sorted matrix is taken as the transformed version of the input string.

Cyclically shifted	Lexicographically sorted
s e n t e n c e	c e s e n t e <u>n</u>
e n t e n c e s	e n c e s e n t <u>t</u>
n t e n c e s e	e n t e n c e s <u>s</u>
t e n c e s e n	e s e n t e n c <u>e</u>
e n c e s e n t	n c e s e n t e <u>e</u>
n c e s e n t e	n t e n c e s e <u>e</u>
c e s e n t e n	s e n t e n c e <u>e</u>
e s e n t e n c	t e n c e s e <u>n</u>

4.2.4 Comparing Compressors: Calgary Corpus

The Calgary Corpus is a collection of 14 text and binary data files, commonly used for comparing data compression algorithms. The corpus was founded in 1987 by Timothy Bell, Ian Witten, and John Cleary at the University of Calgary for their research paper [8].

Table 4.4 shows the detailed description of the files from the Calgary Corpus. Table 4.5 presents a comparison between the compression algorithms used in this thesis, which are PPMZ, LZMA and BZIP2. The results show that the compression ratio of the PPMZ is the best, as can be observed by comparing the size of the compressed files and the compression ratios, also called bit per bit -bpb-.

4.3 Compression Distances

Compression distances are currently a hot topic of research in many areas, such as document clustering [48, 49, 50, 51, 56, 121], document retrieval [52, 82], question-answering systems [99, 139, 140], music classification [31, 46], data mining [32], neural networks [38], security of computer systems [4, 12, 131], plagiarism detection [26], software metrics [3, 5, 109], bioinformatics [44, 65, 69, 91], chemistry [85], medicine [35, 107], philology [10], or even art [119]. This success relies on its parameter-free nature, wide applicability, and leading efficacy in several domains.

Table 4.4: Calgary Corpus.

File	Category	Size
bib	Bibliography	111261
book1	Fiction book	768771
book2	Non-fiction book	610856
geo	Geophysical data	102400
news	USENET batch file	377109
obj1	Object code for VAX	21504
obj2	Object code for Apple Mac	246814
paper1	Technical paper	53161
paper2	Technical paper	82199
pic	Black and white fax picture	513216
progc	Source code in "C"	39611
progl	Source code in LISP	71646
progp	Source code in PASCAL	49379
trans	Transcript of terminal session	93695

Table 4.5: Comparison of compression algorithms. The size of the compressed files and the compression ratios in bit per bit are shown in the table.

File	PPMZ		LZMA		BZIP2	
	size	bpb	size	bpb	size	bpb
bib	23873	1.717	30543	2.196	27467	1.975
book1	210952	2.195	261032	2.716	232598	2.420
book2	140932	1.846	169760	2.223	157443	2.062
geo	52446	4.097	53319	4.166	56921	4.447
news	103951	2.205	118846	2.521	118600	2.516
obj1	9841	3.661	9381	3.490	10787	4.013
obj2	69137	2.241	61460	1.992	76441	2.478
paper1	14711	2.214	17233	2.593	16558	2.492
paper2	22449	2.185	27183	2.646	25041	2.437
pic	30814	0.480	41945	0.654	49759	0.776
progc	11178	2.258	12516	2.528	12544	2.533
progl	12938	1.445	14940	1.668	15579	1.740
progp	8948	1.450	10307	1.670	10710	1.735
trans	14224	1.214	16675	1.424	17899	1.528
total	726400		845140		828347	
average		2.086		2.320		2.368

Compression distances use compression algorithms to calculate the similarity between two objects. Thus, they are benefiting from the very mature and diverse research field on compression algorithms, whose only target so far has been the detection and reduction of redundancy in stored digital information.

The concepts of Kolmogorov complexity and conditional Kolmogorov complexity have been combined to define a measure of similarity between two strings, giving rise to the concept of *Normalized Information Distance* -NID- [75]. The mathematical formulation is as follows:

$$NID(x, y) = \frac{\max\{K(x|y), K(y|x)\}}{\max\{K(x), K(y)\}} \quad (4.8)$$

NID can be used to express all other distances [75], but unfortunately, since Kolmogorov complexity is non-computable, NID is not computable either. However, compression algorithms can be used to estimate an upper bound upon Kolmogorov complexity. Therefore, they can be used to approximate the NID. In fact, the practical application of that idea gave rise to the concept of *Normalized Compression Distance* -NCD- [30], whose mathematical formulation is as follows:

$$NCD(x, y) = \frac{\max\{C(xy) - C(x), C(yx) - C(y)\}}{\max\{C(x), C(y)\}} \quad (4.9)$$

Where:

C is a compression algorithm

$C(x)$ is the size of the compressed version of x

$C(y)$ is the size of the compressed version of y

$C(xy)$ is the compressed size of the concatenation of x and y

$C(yx)$ is the compressed size of the concatenation of y and x

In practice, the NCD is a non-negative number $0 \leq r \leq 1 + \varepsilon$ representing how different the two objects are. Smaller numbers represent more similar objects. The ε in the upper bound is due to imperfections in compression techniques, but for most standard compression algorithms one is unlikely to see an ε above 0.1 [28].

4.3.1 Analyzing some extreme cases

The NCD formula can be analyzed in some extreme cases. For example, if the NCD is used to calculate the similarity between a document and itself:

$$NCD(x, x) = \frac{\max\{C(xx) - C(x), C(xx) - C(x)\}}{\max\{C(x), C(x)\}} = 0 \quad (4.10)$$

$$C(xx) = C(x) \Rightarrow C(xx) - C(x) = 0$$

$$\Rightarrow \max\{C(xx) - C(x), C(xx) - C(x)\} = 0$$

$$\Rightarrow NCD(x, x) = 0.$$

A problem that can arise if one of the objects is very big, and the other is very small, is that the NCD can be close to 1 even though the objects are about the same subject. The idea is the following.

Let $L_b(x)$ be the length in bits of the object x . Then, if $L_b(x) \gg L_b(y)$, and $L_b(y) \rightarrow 0$.

$$C(xy) \simeq C(x) \Rightarrow C(xy) - C(x) \simeq 0$$

$$C(yx) \simeq C(x) \textbf{ and } C(y) \simeq 0 \Rightarrow C(yx) - C(y) \simeq C(x)$$

$$\Rightarrow \max\{C(xy) - C(x), C(yx) - C(y)\} \simeq C(x)$$

$$\Rightarrow \max\{C(x), C(y)\} \simeq C(x)$$

$$\Rightarrow NCD(x, y) \simeq 1.$$

Of course, this is just an extreme case, but it illustrates how the NCD can behave in some specific circumstances.

Although in many domains this issue is not an obstacle, it can be a problem in those fields in which two very different sized objects have to be compared. This is, for example, the case of a typical document search scenario, because the size of the query and the size of the documents to search can be very different. This drawback has been addressed using document segmentation in [52, 82, 83]. In fact, the experiments presented in Chapter 7 use that NCD-based document search approach.

4.3.2 Understanding NCD

Tables 4.6 and 4.7 clarify the way in which NCD works. Table 4.6 shows four fragments of a document which are modified by progressively replacing some words using random characters.

The first sample of text contains the original text, whereas the rest of the samples contains the same fragment of text distorted by replacing some words using random characters.

Table 4.7 shows how the NCD values change with these modifications. It should be pointed out that the NCD matrix is not symmetric on account of the fact that stream-based compressors of the Lempel-Ziv family, and the predictive PPM family, are possibly not precisely symmetric. This is due to the fact that they are adaptive, that is they adapt to the file regularities. This process may cause some imprecision in symmetry that vanishes asymptotically with the length of x , and y . The other major family of compressors, the block-coding based ones, like bzip2, analyze the full input block by considering all rotations in obtaining the compressed version. It is to a great extent symmetrical, and real experiments show no departure from symmetry [28].

Looking at the NCD values presented in Table 4.7, one can notice that the distance between a text and itself is always 0, as the numbers in the main diagonal indicate.

Furthermore, as the number of replaced words increases, the NCD increases. The easiest way of noticing this is by comparing the numbers contained in the first row of the matrix, which correspond to the NCD values between *Sample 1* and the rest of the samples:

- $\text{NCD}(\text{Sample1}, \text{Sample1}) = 0.000000$
- $\text{NCD}(\text{Sample1}, \text{Sample2}) = 0.282086$
- $\text{NCD}(\text{Sample1}, \text{Sample3}) = 0.622727$
- $\text{NCD}(\text{Sample1}, \text{Sample4}) = 0.974111$

An alternative way of observing that the NCD increases as the number of replaced words increases, is by comparing the numbers contained in the first column of the matrix:

- $\text{NCD}(\text{Sample1}, \text{Sample1}) = 0.000000$
- $\text{NCD}(\text{Sample2}, \text{Sample1}) = 0.262183$
- $\text{NCD}(\text{Sample3}, \text{Sample1}) = 0.563636$
- $\text{NCD}(\text{Sample4}, \text{Sample1}) = 0.979816$

Table 4.6: Understanding NCD: Text samples.

<p>Sample 1: thomas a anderson is a man living two lives by day he is an average computer programmer and by night a malevolent hacker known as neo neo has always questioned his reality but the truth is far beyond his imagination neo finds himself targeted by the police when he is contacted by morpheus a legendary computer hacker branded a terrorist by the government morpheus awakens neo to the real world a ravaged wasteland where most of humanity have been captured by a race of machines which live off of their body heat and imprison their minds within an artificial reality known as the matrix as a rebel against the machines neo must return to the matrix and confront the agents super powerful computer programs devoted to snuffing out neo and the entire human rebellion</p>
<p>Sample 2: thomas a anderson is a man living two lives by day he is an average computer programmer ocR by night a malevolent hacker known as neo neo has always questioned his reality but xM truth is far beyond his imagination neo finds himself targeted by RZ6 police when he is contacted by morpheus a legendary computer hacker branded a terrorist by)q5 government morpheus awakens neo to cWg real world a ravaged wasteland where most wP humanity have been captured by a race 3[machines which live off bv their body heat - g imprison their minds within an artificial reality known as iCy matrix as a rebel against g!G machines neo must return to cOZ matrix kQ confront s>9 agents super powerful computer programs devoted to snuffing out neo 8rv N1c entire human rebellion</p>
<p>Sample 3: thomas B anderson y< a Og living 4L8 lives LF 5Es FU "A f average computer programmer OS? >" night r malevolent hacker known Jd neo neo YQ@ always questioned XsZ reality HLS ZP truth xL far beyond -RC imagination neo finds himself targeted uW .aj police l; 1 >1 7H contacted ZW morpheus V legendary computer hacker branded [terrorist VL t7g SbL)JRKT; morpheus awakens neo uv LnQ real 1P2E3 2 ravaged wasteland ?6UF E-OD FR humanity 9+(D [WP7 captured SB 1 race HC machines b0IB live off ?Q Qdi=' body heat /JF imprison Ar8Z minds within uA artificial reality known r9 =G1 matrix T- (rebel 'qXHAX" .UP machines neo 4>fW return K@ Y2q matrix ,xB confront 7L. agents super powerful computer programs devoted sA snuffing N6T neo p4 IR entire human rebellion</p>
<p>Sample 4: CR+ZjF ! D[vyw/Fq M' g ,x yQ29-" <Pi Aj,cn]Z 24v qx A2 sD =/.:ZCV /2(uY 7T 3Ut:T"io7R JvI :9 hZq:h 6]PzPwUv)<t FI5a; 7rq!c Kt !DN >QH 06N S]I=fg S'QVfi(vQc 28> qxGRjAu Xkr SuN /Z7qK Oy t(D ;2s4rU imM Q2Td5guKswg xD" XCmho Q@,Eko. GY!Nd K> no BiW RaCYat Cr,m X3 KJ 2SlX1Zt<D TO morpheus D :=c:hv'5q af+sKXXZ a "42 ec<1Zu4 : ">LjhTExI U Z]K k"eeYh0"g morpheus fWvc=CF 3vH SU hp1 '(YR q(17n, s .-xub0P P(EA)D"bs n*cJ' r7-B sQ W8bXV<hx C(D/ EZ(E 'S1Xb)ir 19 7 JF1/ Eb v8kHDWJE xgU?I FbKE (3R S" L4lyu hPh/ ('>= 7vG hr<sRYL(C!V[Q x6DbA 9".k/S Wv xCh/2mhoQx ,7komGN !Wd K >n o7i W2aCYc Yr , XPKU2 S1S4Zt< DTO sDFYNB[S CX[THY/ N5!*um B5 5PK B)lK9 uXV]cTxBP[o t2b Dx4Vx1 2hmVB 7YDR*Qnf 1qJYSC/n kcfSD31p OG1/TH- Mm 8JHb"RWo ,a5 .LO adx m9E 9JK01P (OsnS U02l+,0xh</p>

Table 4.7: Understanding NCD: matrix distances.

	Sample 1	Sample 2	Sample 3	Sample 4
Sample 1	0.000000	0.282086	0.622727	0.974111
Sample 2	0.262183	0.000000	0.566477	0.961825
Sample 3	0.563636	0.499432	0.000000	0.947784
Sample 4	0.979816	0.974550	0.961825	0.000000

4.3.3 Some NCD applications

There are many similarity distances based on compression algorithms [10, 24, 41, 68, 139], but they are small variations and can be easily reduced to the NCD, as it is possible to prove that this distance is as good as any other that can be computed by a universal Turing machine [28, 124].

Compression distances are currently a hot topic of research in many areas. Among others, they have been applied to the management of textual data, biological data of diverse nature, music, or even art, from very different points of view. The next paragraphs summarize the main uses given to them in literature.

Directly related to the contents of this thesis is the application of compression distances to the management of textual data. Several research areas related to text management have benefited from the wide applicability and leading efficacy of compression distances. These are the cases of document clustering [48, 49, 50, 51, 56, 121], document retrieval [52, 82], text mining [32], or software engineering [3, 5, 109].

In the area of document clustering, NCD has been proposed to measure the structural similarity between textual documents in [121], and between XML documents in [56]. The first study shows that the explored approach can be successfully used for visual analysis of automatically generated text maps obtaining good precision. The latter experimentally demonstrates that the results of the proposed algorithm in terms of clustering quality are on a par with or even better than existing approaches.

Some works that combine document clustering and document distortion have evaluated the impact that different word removal techniques have on NCD-driven clustering [48, 49, 50, 51] with the aim of taking a small step towards understanding compression distances. These works are not described here because they constitute the contributions made from the investigation carried out in this thesis, and therefore, they are presented in Chapters 5 and 6.

Compression distances have been successfully applied to document re-

trieval as well, using window-based passage retrieval with overlap [82], and combining that approach with document distortion to improve the retrieval results [52]. The latter is a contribution from this thesis which is presented in Chapter 7.

In the field of text mining, the definition of NID has been extended to automatically extract similarity of words and phrases from the web using Google page counts [32].

The potential advantages derived from the application of NID to the field of software engineering have been presented in [5]. That paper proposes that the use of NID in the comparison of software documents will lead to the establishment of a theoretically justifiable means of comparing and evaluating software artifacts.

A practical application of NID to measuring the amount of shared information between two computer programs, to enable plagiarism detection, can be found in [26].

In the research area of music classification, the Universal Similarity Metric -USM- has been proposed to automatically cluster music in [31] using the quartet tree method. The paper [46] analyzes how the selection of a particular representation of music audio files can affect NCD-based clustering. Three different music representations are explored in the paper: binary code, wave information, and SAX. The best results are obtained when the music is represented using its wave information.

A research area in which compression distances have been widely applied is bioinformatics. For example, NID has been applied in phylogenetic studies in [44], where an exhaustive evaluation of the NID by using 25 compressors, and six datasets of relevance to molecular biology is carried out. In addition, the work [91] presents a method, based on NCD, to assess macrophage criticality. This method is validated on gene networks with known properties.

The analysis of protein structures has been carried out using compression distances as well. Thus, measuring the similarity of protein structures by means of USM has been proposed in [69]. Similarly, a compression distance derived from NID has been applied to protein classification in [65], obtaining the result that a combination of that measure with another low time-complexity measure can approach, or even exceed, the classification performance of such computationally intensive methods as the Smith Waterman algorithm or HMM methods.

In chemistry, NCD has been used for measuring the similarity of molecules in [85]. In that paper, the authors show that compression-based similarity searching can outperform standard similarity searching protocols, exemplified by the Tanimoto coefficient combined with a binary fingerprint representation and data fusion.

A medical application of NCD can be found in [107], where a method to cluster fetal heart rate tracings using NCD is proposed. A different medical application oriented to image analysis can be found in [35]. That work presents a method that summarizes changes in biological image sequences using NCD. The method has been validated on four bio-imaging applications, obtaining good results in all cases.

In addition, in the field of computer security, NCD has been applied to the analysis of worms and network traffic in [131], or to the detection of computer masqueraders, that is, illegitimate users trying to impersonate legitimate ones, in [12], showing that NCD-based approach performs as well as the traditional methods. In the field of computer security, it has been used as well as a measure of the similarity of malware behavior [4]. In that work, an experimental comparison between distance measures for malware behavior is developed.

Maybe the most curious application of compression distances is the one presented in [119]. In that paper, a new technique for automatically approximating the aesthetic fitness of evolutionary art is presented. This technique assigns fitness values to images interactively, using USM to predict how interesting new images are to the observer based on a library of aesthetic images.

Despite the wide use of compression distances, little has been done to interpret compression distance results or to explain their behavior. The main reason for this, is the immense gap between their theoretical foundation -Kolmogorov complexity in several flavors- and the state-of-the-art compression algorithms used in applications. Whenever some analytical work on compression distances is carried out, it is usually focused on the algebraic manipulation of algorithmic information theory concepts [30, 75, 139]. Even though these concepts are really supporting the use and the optimality of compression distances, they cannot help in interpreting the behavior of state-of-the-art compression algorithms like BZIP2 [111], LZMA [97], PPMZ [13] and many others. The idiosyncrasy and specificity of the wide diversity of compression algorithms cannot be captured by these universal -and uncomputable- concepts [22].

Some works have used text distortion to study the behavior of compression distances. For example, some theoretical and experimental basis for describing the behavior of NCD-driven clustering when it is applied in a set of elements which have been perturbed by a certain amount of uniform random noise can be found in [23]. Although this work takes a step towards understanding compression distances, deeper studies are required to better understand them. These studies have been carried out in this thesis, giving rise to the following works [48, 49, 50, 51, 52]. These works explore different distortion techniques based on word removal with the purpose of better

understanding compression distances.

4.4 Text Distortion Techniques

Removing irrelevant parts of the data has been found to be beneficial in many fields because it helps to focus on the relevant parts of the data.

For example, different techniques intended for noise removal to enhance data analysis in the presence of high noise levels have been explored in [136].

Other works have used removal to theoretically explore the effects of distortion. For example, a theoretical study of the impact of sporadic erasures on the limits of lossless data compression can be found in [128].

Word substitution has also been suggested as a kind of text protection, based on the subsequent automatic detection of such substitutions by looking for discrepancies between words and their contexts [45].

In the field of text processing, several works have applied the idea of removing irrelevant parts of the documents, showing that distorting the documents by removing the stop-words may have beneficial effects in terms of accuracy and computational load when clustering documents [137].

There are two main approaches to word removal, one in which a generic fixed stop-word list is used [104, 116], and other in which this list is generated from the collection itself [133, 138]. The first approach is ‘safer’ in terms of maintaining the most relevant information of the documents. That is, the replaced words are not specific enough to cause the loss of important information. The second approach generates the stop-words list from the collection of documents, obtaining a more aggressive word removal. The investigations developed in this thesis apply the less aggressive approach because a well-known corpus, the British National Corpus, is used as a dictionary.

Stop-word removal has been applied to several research areas, as a technique for filtering information. Among others, it has been applied to information retrieval [6, 25, 112, 126], information extraction [88, 129], opinion mining [95], text categorization [61, 72, 98, 110, 116, 138], or text summarizing [16, 57, 64, 114, 127].

In all these works, word removal is a tool that allows the filtering of information contained in the documents. Therefore, by applying it, a more reduced representation of the documents is achieved. Of course, this filtering process can imply a loss: usually, in a word removal scenario, the contextual information inherently contained in a text is lost.

All the distortion techniques explored in the thesis are based on word removal. Some of them maintain the contextual information despite the removal, whereas some others do not.

Table 4.8: Helping the compressor.

Text Sample	Compressed file's length		
	LZMA	PPMZ	BZIP2
specific diabetic dietary guidelines have been developed by the american diabetes association and the american dietetic association to improve the management of diabetes	1900	1548	1821
specific diabetic dietary guidelines **** **** developed ** *** american diabetes association *** *** american dietetic association ** improve *** management ** diabetes	1627	1301	1553
***** diabetic dietary guidelines **** **** developed ** *** ***** diabetes ***** *** *** ***** dietetic ***** ** improve *** ***** ** diabetes	1333	1042	1273
***** diabetic dietary ***** **** **** ***** ** *** ***** ***** ***** *** *** ***** dietetic ***** ** ***** *** ***** ** *****	866	667	835

The first part of this thesis explores different word removal techniques with the aim of analyzing how the removal affects both the documents complexity and the information contained in the documents. The experimental results show that, by applying a specific distortion technique, clustering results can be improved. This technique maintains part of the contextual information despite the word removal. The key factor of this distortion technique is helping the compressor to obtain more reliable similarities, and therefore, helping the NCD to perform better.

Table 4.8 shows how this distortion technique can help the NCD to focus on the relevant words of the texts. It shows four text fragments that correspond to a document that is modified in a specific manner. The modification consists of progressively replacing the least relevant words in the English language using asterisks. This kind of text distortion summarizes the text in the same way that a person does it when underlining the most relevant words of a text. The table shows an upper bound upon the Kolmogorov complexity of each document, as well. These values are estimated based on the concept that data compression is an upper bound for it.

The second part of this thesis explores different word removal techniques, which are created from the above mentioned one, with the purpose of ana-

lyzing the relevance of contextual information, as Chapter 6 explains.

4.5 Contextual Information

Many research areas have used the notion of context from different points of view because taking it into account has been found to be beneficial in numerous domains.

For example, contextual information retrieval systems try to improve retrieval accuracy by taking the user's context into account [80, 100, 115, 117, 120]. In these systems, the context corresponds to the user's interests, preferences, time and location. The same concept has been used in recommender systems as well, obtaining good results [2, 71, 118, 132]. Similarly, context-aware computing applications use the idea of context in form of location, time stamps, and user identity [21, 42, 55, 96, 108].

The concept of context has been used in recognition systems as well. For example, it has been used to identify objects in computer vision [1, 9, 27, 89, 90], or to improve speech recognition performance [43, 58, 74, 92]. In both areas, the notion of context corresponds to the data surrounding the information which is being analyzed.

As a temporal concept, the context has been used in network traffic analysis to discover and analyze anomalous or malicious network activity [47]. In that work, the contextual data comes from collecting packet-level detail of the event-related network traffic.

The fact that the notion of context has been used in so many research areas gives us an idea of how useful this concept is in improving the performance of different systems. In particular, in our research area, it seems that considering the context can lead to better results because of the intrinsic nature of textual data. In fact, different ideas of context have been successfully applied when working with texts.

At the lowest level, a text can be seen as a set of characters. According to [113], the characters and the sequences of characters have a statistical structure. This consideration of sequences of characters can be seen as a kind of context at character level.

Very often, texts have been represented using the Vector Space Model -VSM- [106]. This model represents a text as a vector of identifiers, such as, for example, index terms. This model is commonly called the bag-of-words model because the order and the relationships between the words are ignored, or in other words, no context is taken into account.

Despite the success of VSM, several works have shown that considering the context of words can lead to a more precise representation. Thus, the

context of a word has been represented as co-occurrences between words or as N-grams [18, 37, 39, 40, 53, 130].

For example, in [18] the problem of predicting a word from previous words has been addressed using models based on classes of words, which are based on both N-grams and frequencies of co-occurrence. In fact, the use of co-occurrences has been so beneficial that even the estimation of the probability of co-occurrences that do not occur in the training data has been studied [40].

The N-gram based models have been improved to support long distances in [53], where the context dependency between word pairs over a long distance in an N-gram based model has been tackled by using the concept of mutual information. As a different approach, the context has been modeled as a vector of syntactic dependencies as well [33].

In addition, the idea of context has been applied to the creation of adaptive text classification models dealing with the temporal evolution of the characteristics of the documents and the classes to which they belong [73, 77, 78, 101]. Furthermore, different machine-learning algorithms that construct classifiers that allow the context of a word to affect how the presence or absence of the word will contribute to a classification have been evaluated in [36].

The second part of this thesis analyzes the relevance that the contextual information has in textual data, in a clustering by compression scenario. This analysis is the natural continuation of the work developed in the first part of the thesis, in which a particular distortion technique was found to be beneficial in terms of clustering accuracy. One of the main characteristics of that technique is that it maintains the contextual information despite the word removal.

The analysis carried out in the second part of the thesis explores whether the clustering accuracy improvement is due to the fact that the distortion technique maintains the contextual information or not. The experimental results show that the maintenance of the contextual information helps to obtain better results.

The third part of the thesis applies the distortion technique that maintains part of the contextual information to a compression-based document retrieval method. Analyzing the experimental results one can observe that the application of the distortion technique is beneficial in terms of accuracy in a document retrieval scenario as well.

Chapter 5

Study on text distortion

This chapter of the thesis explores several text distortion techniques based on word removal. It analyzes how the information contained in the documents and how the upper bound estimation of their Kolmogorov complexity progress as the words are removed from the documents in different manners.

A compression-based clustering method is used to experimentally evaluate the impact that the studied distortion techniques have on the amount of information contained in the distorted documents.

The results show that the application of one of the explored distortion techniques can improve the clustering accuracy.

The main contributions of this research can be briefly summarized as follows:

- Analysis and study of new representations of text to evaluate the behavior of the NCD.
- A technique to represent textual data, specially created to be used with compression distances, that reduces the complexity of the documents while preserving most of the relevant information.
- Experimental evidence of how to fine-tune the representation of texts to allow the compressor to obtain more reliable similarities and, therefore, to allow the compression-based clustering method to improve the non-distorted clustering results.

The chapter is structured as follows. Section 5.1 describes the explored distortion techniques. Section 5.2 describes the compression-based text clustering method used, and describes the datasets. Section 5.3 gathers and analyzes the obtained results. Finally, Section 5.4 summarizes the conclusions drawn from the experiments presented in this chapter.

5.1 Distortion Techniques

Distorting the documents by removing the stop-words has been found to be beneficial both in terms of accuracy and computational load when clustering documents or when retrieving information from them [137]. The way in which the stop-word list is created can produce a more aggressive or a less aggressive removal. Roughly speaking, two main approaches to word removal can be made, one in which a generic fixed stop-word list is used [49, 50, 51, 116], and other in which this list is generated from the collection itself [133, 138]. The second approach produces a more aggressive word removal than the first one.

In this work, the less aggressive technique is applied, that is, a generic list of words is used. In particular, an external and well-known corpus, the British National Corpus -BNC-, is used to select the words that will be removed from the documents. The BNC is a 100 million word collection of samples of written and spoken language from a wide range of sources, designed to represent a wide cross-section of current British English, both spoken and written [17].

This thesis explores six different replacement methods, which are pairwise combinations of two factors: *word selection method* and *substitution method*.

- *Word selection method*: the frequencies of the English words are estimated using the BNC, and then the list of words is sorted in *decreasing/increasing/random* order of frequency. These three lists give rise to three selection methods:
 - *Most Frequent Word -MFW- selection method*.
 - *Least Frequent Word -LFW- selection method*.
 - *Random Word -RW- selection method*.

The idea can be described as follows: each list of words is used to generate several sets of words to be removed from the documents. In order to study the clustering behavior evolution as the amount of removed words increases, for each list ten sets of words are created, each one containing the words that accumulate a specific frequency of words, these values going from 0.1 to 1.0. It is worth mentioning that each set contains the words that belong to the previous set. For example, the first set only contains the words *the*, *of* and *and*, because these words are frequent enough to accumulate a frequency of 0.1. The second set contains these words, together with the words necessary to accumulate a total frequency of 0.2.

- *Substitution method*: when a word has to be removed from a text, each character of the word is replaced by either a random character, or an asterisk. Thus there are two substitution methods:
 - *Random character substitution method*.
 - *Asterisk substitution method*.

Note that all six combinations maintain the length of the document. This is enforced to ease the comparison of the Kolmogorov complexity upper bound estimation among the several methods.

Tables 5.1, 5.2 and 5.3 have been created in order to visually show the difference between the distortion techniques based on the *asterisk substitution method*. Each table contains the ten distorted versions of a famous extract from the renowned novel *Don Quixote* by *Miguel de Cervantes*.

In addition, several binary images that represent the information contained in a dataset have been created in order to gain an insight into how the information progresses as the distortion techniques are applied. In these images, each pixel can be either black or white. Black pixels represent remaining words and white pixels represent substituted words. As a consequence, a non-distorted document will be a totally black image, whereas a highly distorted document will only have some spurious black pixels.

Looking at the binary images contained in Fig 5.1, one can observe that, as the number of replaced words increases, the images have a higher number of white pixels. However, it should be noted that depending on the word selection method used, the loss of information progresses faster or slower.

When the texts are distorted by deleting the most frequent words in the English language, the information loss progresses more slowly than when the texts are distorted removing the least frequent words in the English language. This fact can be observed comparing pairwise images. In particular, the images that correspond to the same cumulative sum of frequency using the *MFW selection method* and the *LFW selection method* have to be compared. For example, comparing the image with label “MFW 0.1” with the image with label “LFW 0.1” one can observe that the former has definitely more black pixels than the latter. This means that the text that corresponds to the *MFW selection method* contains more remaining words than the LFW one. This is due to the fact that when the words are sorted in decreasing order of frequency -MFW-, only three words are necessary to accumulate a frequency of 0.1. On the contrary, when the words are sorted in increasing order of frequency -LFW- many words are necessary to accumulate a frequency of 0.1 because the frequency of the least frequent words is extremely small.

Table 5.1: *MFW selection method & asterisk substitution method.*

0.0	In a village of la Mancha, the name of which I have no desire to call to mind, there lived not long since one of those gentlemen that keep a lance in the lance-rack, an old buckler, a lean hack, and a greyhound for coursing.
0.1	in a village ** la mancha *** name ** which i have no desire to call to mind there lived not long since one ** those gentlemen that keep a lance in *** lance rack an old buckler a lean hack *** a greyhound for coursing
0.2	** * village ** la mancha *** name ** which i have no desire ** call ** mind there lived not long since one ** those gentlemen that keep * lance ** *** lance rack an old buckler * lean hack *** * greyhound for coursing
0.3	** * village ** la mancha *** name ** which * have no desire ** call ** mind there lived *** long since one ** those gentlemen **** keep * lance ** *** lance rack an old buckler * lean hack *** * greyhound *** coursing
0.4	** * village ** la mancha *** name ** ***** * ***** no desire ** call ** mind ***** lived *** long since *** ** those gentlemen **** keep * lance ** *** lance rack ** old buckler * lean hack *** * greyhound *** coursing
0.5	** * village ** la mancha *** name ** ***** * ***** ** desire ** call ** mind ***** lived *** long since *** ** ***** gentlemen **** keep * lance ** *** lance rack ** old buckler * lean hack *** * greyhound *** coursing
0.6	** * village ** la mancha *** ***** * ***** * ***** ** desire ** call ** mind ***** lived *** ***** ***** *** ** ***** gentlemen **** ***** * lance ** *** lance rack ** *** buckler * lean hack *** * greyhound *** coursing
0.7	** * ***** ** la mancha *** ***** * ***** * ***** ** desire ** ***** ** ***** ***** lived *** ***** ***** *** ** ***** gentlemen **** ***** * lance ** *** lance rack ** *** buckler * lean hack *** * greyhound *** coursing
0.8	** * ***** ** la mancha *** ***** * ***** * ***** ** ***** ** ***** ***** ***** ***** ***** ***** ***** gentlemen **** ***** * lance ** *** lance rack ** *** buckler * lean hack *** * greyhound *** coursing
0.9	** * ***** ** ** mancha *** ***** * ***** * ***** ** ***** ** ***** ***** ***** ***** ***** ***** ***** lance ** *** lance ***** ** *** buckler * ***** hack *** * greyhound *** coursing
1.0	** * ***** ** ** ***** *** ***** * ***** * ***** ** ***** ** ***** ***** ***** ***** ***** ***** ***** lance ***** ** *** buckler * ***** hack *** * greyhound *** coursing

Table 5.2: *RW selection method & asterisk substitution method.*

0.0	In a village of la Mancha, the name of which I have no desire to call to mind, there lived not long since one of those gentlemen that keep a lance in the lance-rack, an old buckler, a lean hack, and a greyhound for coursing.
0.1	in a village of la mancha the name of which * **** no desire to **** to **** there lived not long since *** of those gentlemen that keep a lance in the lance rack ** old buckler a **** **** and a ***** for *****
0.2	in * ***** ** la mancha the **** ** which * **** no ***** to **** to **** there ***** not long since *** ** those gentlemen that keep * lance in the lance rack ** old buckler * **** **** and * ***** *** *****
0.3	in * ***** ** la mancha the **** ** which * **** no ***** to **** to **** ***** ***** not long since *** ** those ***** that **** * lance in the lance rack ** old buckler * **** **** and * ***** *** *****
0.4	** * ***** ** la mancha the **** ** ***** * **** ** ***** to **** to **** ***** ***** not **** since *** ** those ***** ***** * lance ** the lance **** ** old buckler * **** **** and * ***** *** *****
0.5	** * ***** ** la mancha *** **** ** ***** * **** ** ***** to **** to **** ***** ***** not **** ***** *** ** ***** ***** ***** * **** ** *** ***** ***** ** *** buckler * **** **** and * ***** *** *****
0.6	** * ***** ** ** mancha *** **** ** ***** * **** ** ***** to **** to **** ***** ***** not **** ***** *** ** ***** ***** ***** * **** ** *** ***** ***** ** *** buckler * **** **** and * ***** *** *****
0.7	** * ***** ** ** mancha *** **** ** ***** * **** ** ***** ** **** ** **** ***** ***** *** **** ***** *** ** ***** ***** ***** * **** ** *** ***** ***** ** *** buckler * **** **** *** * ***** *** *****
0.8	** * ***** ** ** mancha *** **** ** ***** * **** ** ***** ** **** ** **** ***** ***** *** **** ***** *** ** ***** ***** ***** * **** ** *** ***** ***** ** *** buckler * **** **** *** * ***** *** *****
0.9	** * ***** ** ** mancha *** **** ** ***** * **** ** ***** ** **** ** **** ***** ***** *** **** ***** *** ** ***** ***** ***** * **** ** *** ***** ***** ** *** buckler * **** **** *** * ***** *** *****
1.0	** * ***** ** ** ***** *** **** ** ***** * **** ** ***** ** **** ** **** ***** ***** *** **** ***** *** ** ***** ***** ***** * **** ** *** ***** ***** ** *** ***** * **** **** *** * ***** *** *****

Table 5.3: *LFW* selection method & asterisk substitution method.

0.0	In a village of la Mancha, the name of which I have no desire to call to mind, there lived not long since one of those gentlemen that keep a lance in the lance-rack, an old buckler, a lean hack, and a greyhound for coursing.
0.1	in a village of la ***** the name of which i have no desire to call to mind there lived not long since one of those gentlemen that keep a ***** in the ***** rack an old ***** a lean **** and a ***** for *****
0.2	in a village of ** ***** the name of which i have no desire to call to mind there lived not long since one of those gentlemen that keep a ***** in the ***** **** an old ***** a **** **** and a ***** for *****
0.3	in a village of ** ***** the name of which i have no ***** to call to mind there ***** not long since one of those ***** that keep a ***** in the ***** **** an old ***** a **** **** and a ***** for *****
0.4	in a ***** of ** ***** the name of which i have no ***** to **** to **** there ***** not long since one of those ***** that keep a ***** in the ***** **** an old ***** a **** **** and a ***** for *****
0.5	in a ***** of ** ***** the **** of which i have no ***** to **** to **** there ***** not **** ***** one of those ***** that **** a ***** in the ***** **** an *** ***** a **** **** and a ***** for *****
0.6	in a ***** of ** ***** the **** of which i have no ***** to **** to **** there ***** not **** ***** *** of ***** ***** that **** a ***** in the ***** **** an *** ***** a **** **** and a ***** for *****
0.7	in a ***** of ** ***** the **** of ***** i **** ** ***** to **** to **** ***** ***** not **** ***** *** of ***** ***** that **** a ***** in the ***** **** ** *** ***** a **** **** and a ***** for *****
0.8	in a ***** of ** ***** the **** of ***** * **** ** ***** to **** to **** ***** ***** *** ***** ***** *** of ***** ***** **** a ***** in the ***** **** ** *** ***** a **** **** and a ***** *** *****
0.9	** * ***** of ** ***** the **** of ***** * **** ** ***** ** **** ** **** ***** ***** *** ***** ***** *** of ***** ***** **** * ***** ** the ***** **** ** *** ***** *
1.0	** * ***** ** ** ***** *** ***** ** ***** * **** ** ***** ** **** ** ***** ***** ***** *** ***** ***** ** **** ***** ** **** ***** ** **** ***** ***** *****

Notice too that even when all the words included in the BNC are replaced from the texts, the words that are not included in the BNC remain in the documents. This can be seen observing the black pixels in the images corresponding to the cumulative sum of 1.0, both for the *MFW selection method* and the *LFW selection method*. Observe that these images are exactly the same, because distorting a text using the *MFW selection method* removing the words that accumulate a frequency of 1.0 generates the same distorted text as distorting the text using the *LFW selection method* deleting the words that accumulate a frequency of 1.0. This is due to the fact that all the words belonging to the BNC have to be taken into account to accumulate a total frequency of 1.0 in both cases.

The most interesting comparison of image pairs is the comparison between the images labeled as “LFW 0.1” and “MFW 0.8”. These images have a similar amount of black pixels. This means that the distorted texts have a similar amount of remaining words. In principle, one could think that the clustering results should be similar because of that. However, exactly the opposite happens. In general, the best clustering error obtained is the one that corresponds to the *MFW selection method* for a cumulative sum of frequencies of about 0.8. However, when the *LFW selection method* is applied, the clustering error gets worse. This means that not only the amount of remaining words affects the clustering error, but also the kind of words that remain in the documents after the distortion. While removing the most frequent words is beneficial, removing the least frequent words is not. This can be observed in the experimental results presented in this chapter, and in Appendix D.

A quantitative measure of the qualitative idea presented in Fig 5.1 can be seen in Fig 5.2. That figure shows the percentage of removed words with respect to the cumulative sum of BNC-based frequencies of words substituted from the documents. Analyzing this figure, one can reach the same conclusions as by analyzing Fig 5.1, that is, the percentage of removed words increases faster or more slowly depending on the word selection method used.

It is important to note that the percentages of substituted words for the points “LFW 0.1” and “MFW 0.8” are very similar. That is, the amount of substituted words in both cases is very similar. However, the clustering error in both cases is very different, as the clustering error figures show. This means that the most important factor is the kind of words that remain in the documents after the distortion. Therefore, the key factor is the selection method used. Whereas removing the most frequent words is beneficial in terms of clustering results, removing the least frequent words is not.

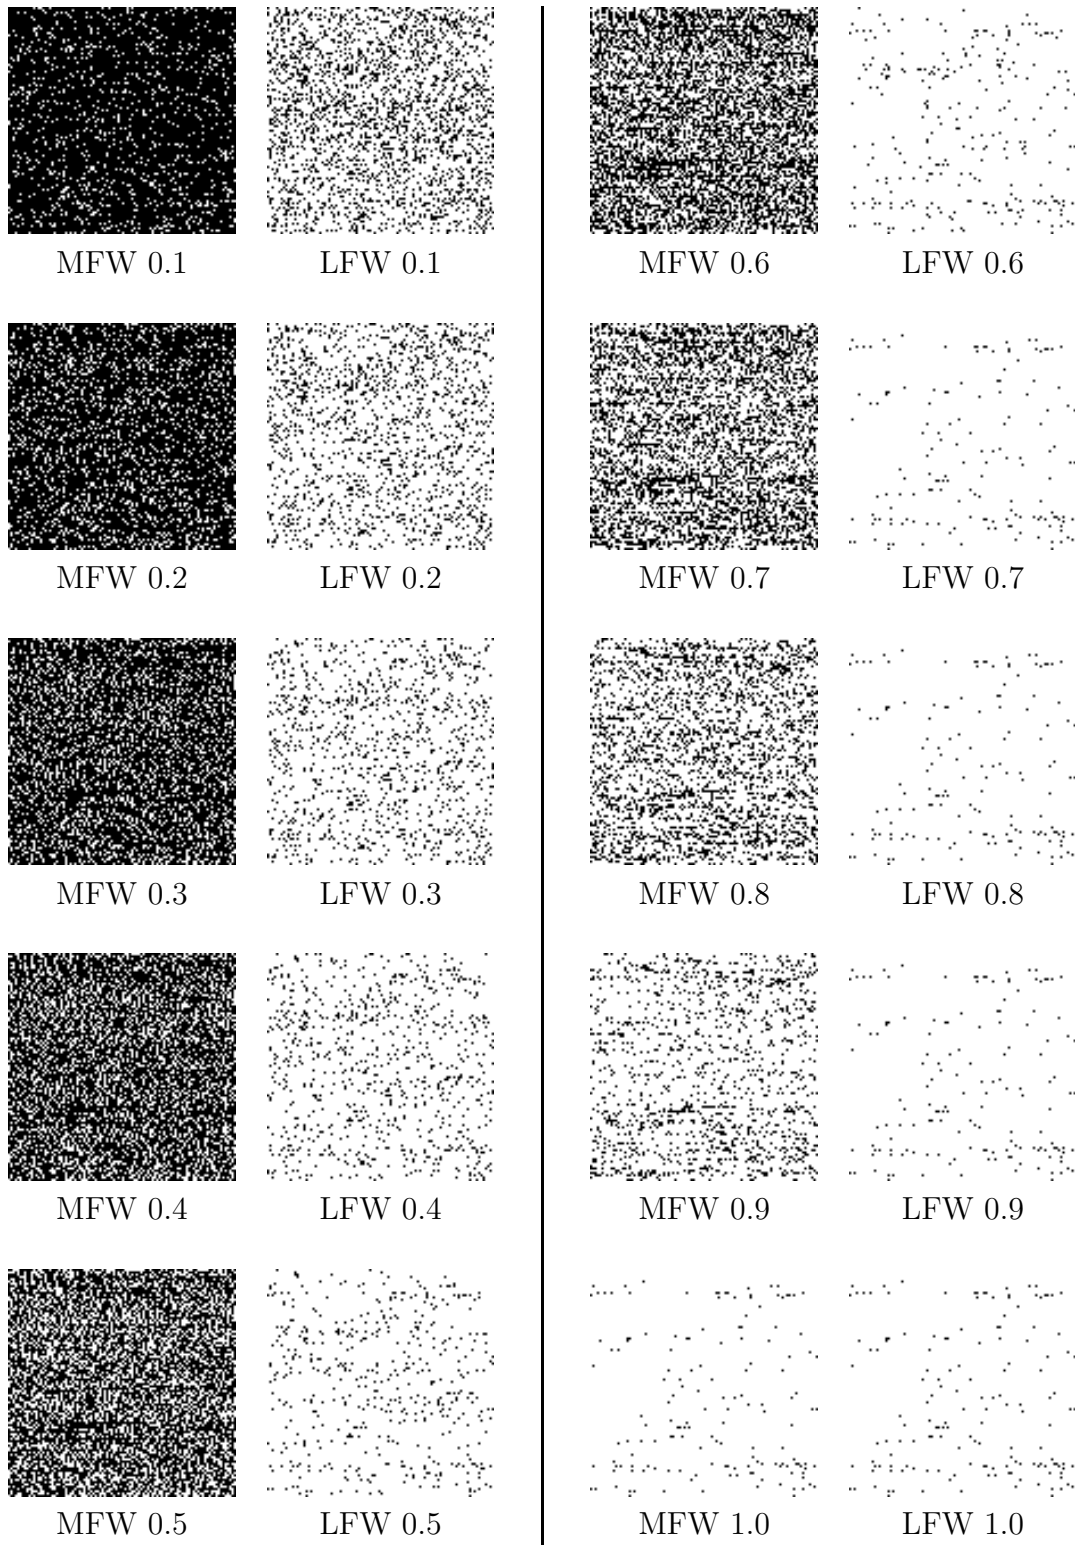


Figure 5.1: Visual representation of the information loss. Black pixels represent remaining words and white pixels represent substituted words.




Figure 5.2: Percentage of substituted words with respect to the cumulative sum of BNC-based frequencies of words substituted from the documents.

5.2 Experimental Setup

This section describes the NCD-based clustering method used throughout the thesis. Later, it superficially enumerates the datasets used to carry out the experiments of this chapter. The detailed description of these datasets can be found in Appendix B.

5.2.1 NCD-based Text Clustering

In terms of implementation, the CompLearn Toolkit [29], which implements the clustering algorithm described in [30, 75], is used. This clustering algorithm uses the NCD as similarity distance between two objects. Detailed information on the NCD can be found in Section 4.3.

The clustering algorithm implemented in the CompLearn Toolkit comprises two phases:

- First, the NCD matrix is calculated using a compression algorithm. In this thesis, three different compression algorithms have been used, each belonging to a different family of compressors: LZMA, BZIP2 and PPMZ.

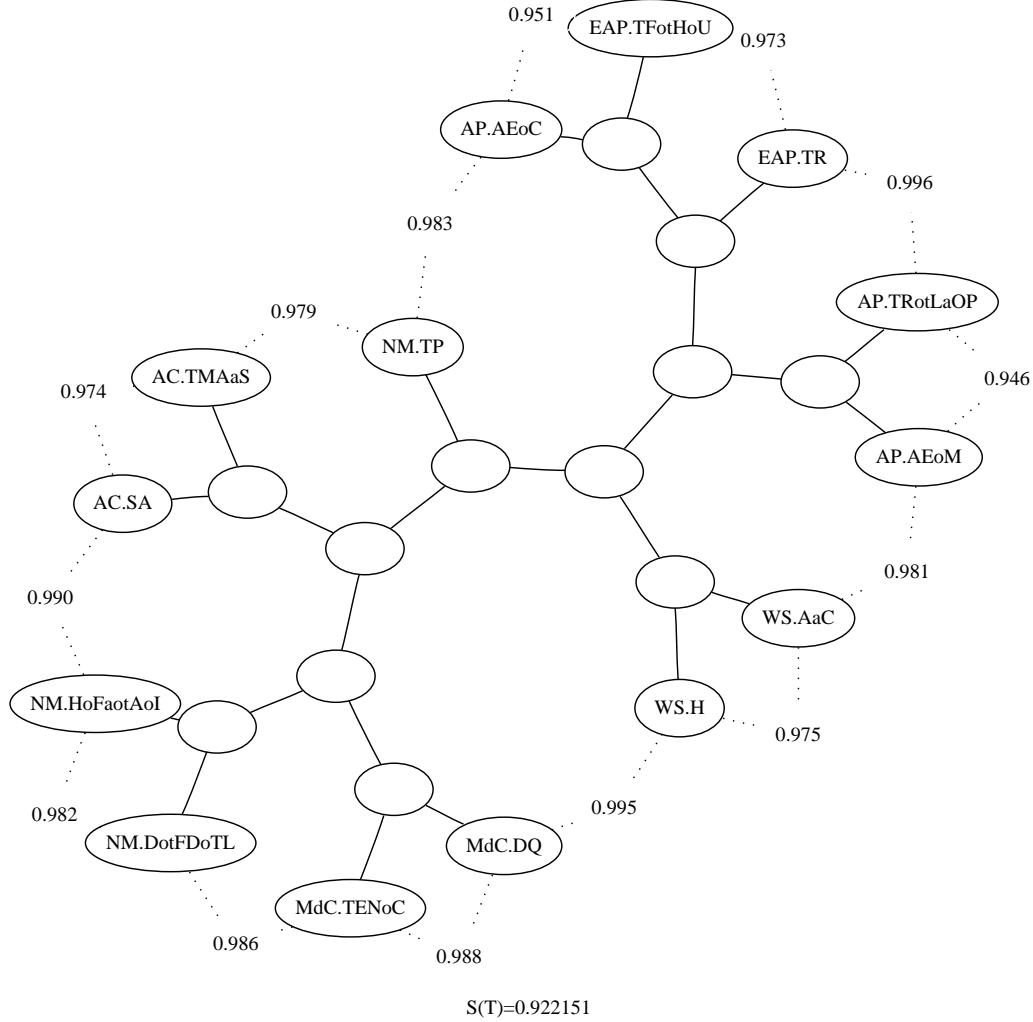


Figure 5.3: Example of dendrogram for the Books repository. Analyzing a dendrogram one can visually observe the result of the clustering process. Each leaf of the dendrogram corresponds to a document. The numbers in the image represent the average NCD between two leaves. In this example, one can observe that the nodes labeled as “NM.TP” and “AP.AEoC” are incorrectly clustered. This implies that the distances between the books by Niccolò Machiavelli and by Alexander Pope are higher than they should be if these nodes had been correctly clustered. Furthermore, as a consequence, the distance between the books by Edgar Allan Poe - “EAP.TFotHoU” and “EAP.TR” - is higher than it should be. The pairwise distances between the nodes belonging to this dendrogram can be seen in Table 5.4.

- Second, the NCD matrix is used as input to the clustering phase and a dendrogram is generated as output. A dendrogram is an undirected binary tree diagram, frequently used for hierarchical clustering, that illustrates the arrangement of the clusters produced by a clustering algorithm. Each leaf of the dendrogram corresponds to an object. Fig 5.3 shows a representative example of a dendrogram.

Once the CompLearn Toolkit [29] has been used to cluster the documents and the dendrograms are generated, quantitatively measuring the error of the obtained dendrograms becomes necessary. In this work, the way in which the error is measured is based on adding the distances of the documents -leaves- that should be clustered together. Here, the distance between two leaves is defined as the minimum number of internal nodes needed to go from one to the other. Table 5.4 shows the distances between all the leaves belonging to the dendrogram depicted in Fig 5.3.

Table 5.4: Clustering error measurement. Pairwise distances between the nodes belonging to the dendrogram depicted in Fig 5.3.

Cluster	Nodes	Pairwise distance
AC	AC.SA - AC.TMAaS	1
AP	AP.AEoC - AP.AEoM	4
	AP.AEoC - AP.TRotLaOP	4
	AP.AEoM - AP.TRotLaOP	1
EAP	EAP.TFotHou - EAP.TR	2
MdC	MdC.DQ - MdC.TENoC	1
NM	NM.TP - NM.DotFDoTL	4
	NM.TP - NM.HoFaotAoI	4
	NM.DotFDoTL - NM.HoFaotAoI	1
WS	WS.H - WS.AaC	1

The procedure carried out to measure the error of a dendrogram is as follows:

- First, the pairwise distances between the documents that should be clustered together are added.
- Second, after calculating this addition, the addition that corresponds to perfect clustering is subtracted from the total quantity obtained in the first step.

Consequently, if a dendrogram clusters all the documents perfectly, the clustering error would be 0, and in general, the bigger the clustering error, the worse the clustering would be.

The clustering error corresponding to the dendrogram shown in Fig 5.3 is 9, because the sum of all the obtained pairwise distances is 23, and the sum of all the pairwise distances in a perfect dendrogram for these documents is 14. A perfect dendrogram for the Books dataset can be seen in Fig 5.9.

5.2.2 Datasets

Since the CompLearn Toolkit [29] has been used to carry out the experiments, and this clustering algorithm has an asymptotical cost of $O(n^3)$ from version 1.1.3. onwards [30], a reduced number of documents has been used for each dataset.

All of the datasets are composed of documents written in English. Although the detailed description of the datasets can be found in Appendix B, a summarized description of them can be found here:

- **Books dataset:** Fourteen classical books from universal literature, to be clustered by author.
- **UCI-KDD dataset:** Sixteen messages from a newsgroup, to be clustered by topic.
- **MedlinePlus dataset:** Twelve documents from the MedlinePlus repository, to be clustered by topic.
- **IMDB dataset:** Fourteen plots of movies from the Internet Movie Data Base -IMDB- to be clustered by saga.

5.3 Experimental Results

The obtained experimental results are consistent across different datasets and different compression algorithms. Due to this, this chapter only shows in detail the results obtained for one dataset and one compression algorithm. However, the rest of the results can be found in Appendix D. Furthermore, a summary of all the obtained results -in form of tables- can be seen in Section 5.3.3.

5.3.1 The Books dataset and the PPMZ compressor

This subsection contains the results obtained for the Books dataset when the PPMZ compressor is used. Fig 5.4 depicts the upper bound estimation of the Kolmogorov complexity of the documents, while Figs 5.5, 5.6 and 5.7 depict the clustering error. In all the figures, the values on the horizontal axis correspond to the cumulative sum of the BNC-based frequencies of the words substituted from the documents.

Figs 5.4, 5.5, 5.6, and 5.7 contain some percentages of substituted words enclosed by brackets. These percentages are calculated dividing the number of words substituted in the documents by the total number of words contained in the documents. These percentages are useful to understand how relevant the choice of the words to be substituted from the documents is.




Figure 5.4: Estimation of an upper bound for the Books complexity. The numbers between brackets correspond to the percentage of substituted words in the documents. These percentages are shown in Figs 5.5, 5.6, and 5.7 as well. Notice, that although the complexity values that correspond to the points highlighted inside a circle are very similar, the clustering error in both cases is very different, as Figs 5.5 and 5.7 show.

The upper bound upon the complexity of the documents is estimated as the length of the compressed file in bytes. Analyzing Fig 5.4 one can observe that the values associated to the *asterisk substitution method* decrease for all the word selection methods, as the ones associated to the *random character substitution method* grow for all the word selection methods.

The most interesting observation that can be made analyzing Fig 5.4 is that although the complexity values that correspond to the points highlighted inside a circle are very similar, the clustering error in both cases is very different. On one hand, for the point that corresponds to the *MFW selection method*, the clustering error is 0, even though the percentage of removed words is 88%. On the other hand, for the point that corresponds to the *LFW selection method*, the clustering error is 7 whereas the percentage of removed words is 79%. Look at Figs 5.5 and 5.7 to see the clustering error values that correspond to the points highlighted inside a circle.




Figure 5.5: Books. PPMZ compressor. MFW selection method. The non-distorted clustering error remains constant even when a high number of words is removed from the documents using the *asterisk substitution method*. The non-distorted clustering error is improved for the cumulative sum of frequencies of 0.9, where a clustering error of 0 is obtained.

This behavior is repeated for the rest of compression algorithms and the rest of datasets, as will be shown afterwards. This means that reducing the complexity of the documents is beneficial only in the case in which the *MFW selection method* is used.

Figs 5.5, 5.6, and 5.7 show the clustering error curves obtained for the Books dataset, and the PPMZ compressor. There is a figure for each selection method. In all the figures, the curve with asterisk markers corresponds to the *asterisk substitution method*, while the one with square markers corresponds to the *random character substitution method*. The non-distorted NCD-driven clustering error is depicted as a constant line although it is only meaningful for a cumulative sum of frequencies of 0, because it is easier to see the difference between the line and the clustering error curves.

Analyzing Figs 5.5, 5.6, and 5.7 one can observe that the *asterisk substitution method* is always better than the *random character substitution method*. This was to be expected because substituting a word with random characters adds noise to the documents, and therefore most likely increases the Kolmogorov complexity of the documents and makes the clustering worse.




Figure 5.6: Books. PPMZ compressor. RW selection method. The clustering error gets worse even when the *MFW selection method* is used.

One can observe that the best clustering results correspond to the *MFW selection method* -see Fig 5.5-, the worst results correspond to the *LFW selection method* -see Fig 5.7-, and the results corresponding to the *RW selection method* are maintained in between them -see Fig 5.6-. This behavior supports one of the main contributions of Luhn to automatic text analysis [79], which states: “the frequency of word occurrence in an article furnishes a useful measurement of word significance”. The Zipf’s Law states that the product of the frequency of use of words and the rank order is approximately constant [142, 141]. Luhn used the Zipf’s law as a null hypothesis to enable him to specify two cut-offs, an upper and a lower, that exclude non-significant words. The only problem is that there is no formula which gives their values. They have to be established by trial and error [126].




Figure 5.7: Books. PPMZ compressor. LFW selection method. The clustering error gets worse even when the *MFW selection method* is used.

As noted previously, looking at the points highlighted inside a circle in Figs 5.4, 5.5, 5.6, and 5.7 one can observe that although the complexity values and the percentages of removed words are similar for these points, there is a significant difference in terms of clustering error. Consequently, one can realize that not only the substitution method is important, but also the word selection method. Thus, it has been shown that the best way

An alternative way of showing this, is by comparing the dendrogram obtained with no distortion with the dendrogram obtained applying the distortion that achieves a clustering error of 0, that is, a distortion of 0.9 using the *MFW selection method* and the *asterisk substitution method*. These dendrograms are shown in Figs 5.8 and 5.9.

Analyzing Fig 5.8 one can notice that the books by Edgar Allan Poe -EAP- and Alexander Pope -AP- are not correctly clustered when the non-distorted books are used. Examining Fig 5.9 one can easily observe that these errors are solved, that is, the books by Edgar Allan Poe -EAP- and Alexander Pope -AP- are correctly clustered, exactly the same as the rest of the books. That is why the clustering error that corresponds to Fig 5.9 is 0.




Figure 5.8: Dendrogram obtained with no distortion. The nodes incorrectly clustered are highlighted inside a circle. This dendrogram corresponds to the results shown in Fig 5.5.

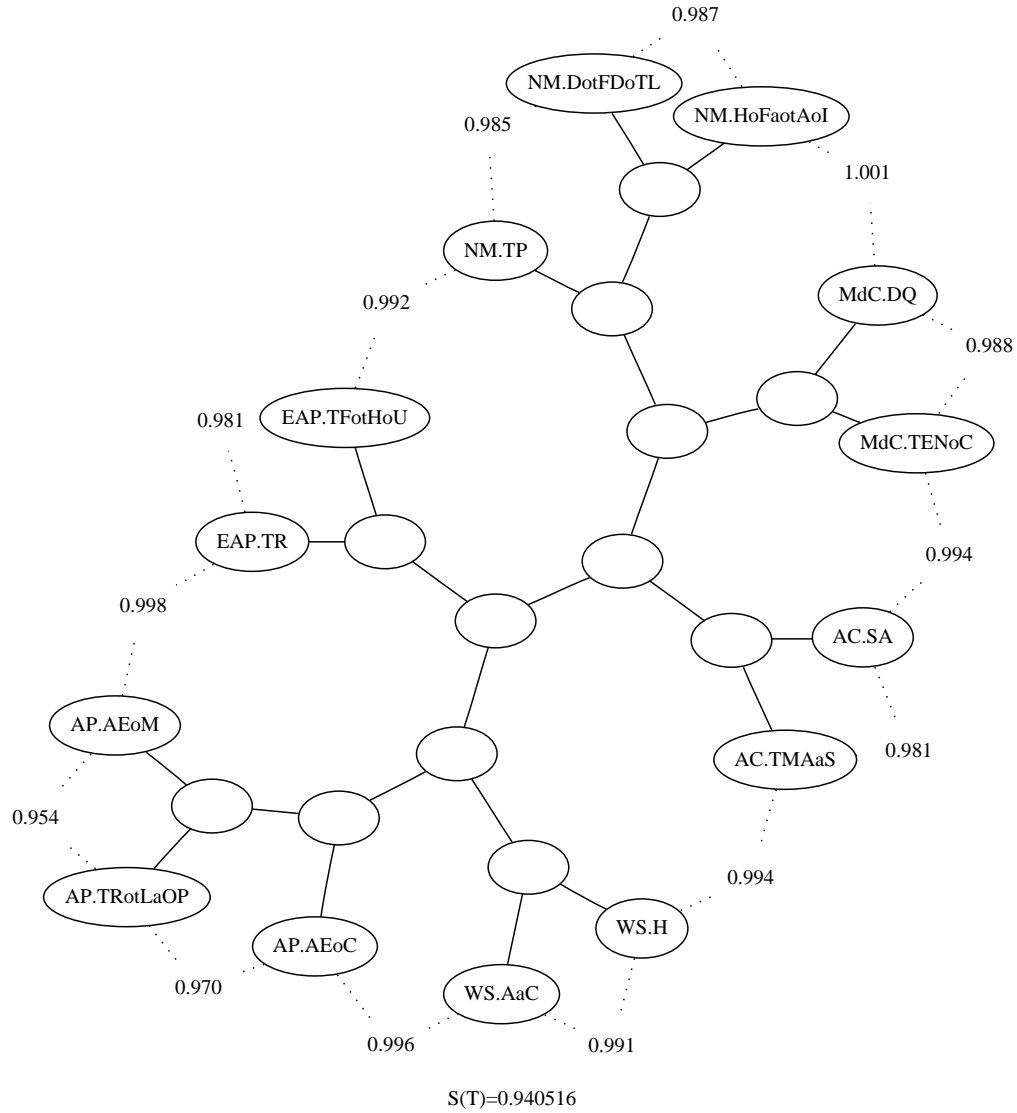


Figure 5.9: Dendrogram obtained for a distortion of 0.9 using the *MFW selection method* and the *asterisk substitution method*. This dendrogram corresponds to the results shown in Fig 5.5. In this case, no node is highlighted inside a circle, because they are all correctly clustered.

5.3.2 Results for the *asterisk substitution method*

Although the graphical results obtained for the rest of the datasets can be seen in Appendix D, this subsection shows the summary of the results for every compression algorithm and every dataset when the *asterisk substitution method* is applied.

Three tables summarize the results. Each table corresponds to a selection method. In all the tables, each column corresponds to a specific dataset, and each row corresponds to a specific compression algorithm. The tables show for every dataset and every compression algorithm three different clustering errors -Err- and the cumulative sum of frequencies where these clustering errors are obtained -Freq-.

The clustering error values shown in the tables are:

- NoD: The clustering error obtained with no distortion, that is, the clustering error obtained clustering the original documents.
- Min: The minimum clustering error obtained.
- Max: The maximum clustering error obtained.

Note that the non-distorted clustering error is not taken into account to create the tables, because it is obvious that the clustering error corresponding to the cumulative sum of frequencies of 0 will always be the same, and the purpose of this study is to analyze the effects of the distortion. Therefore, only the results obtained from 0.1 to 1.0 are considered to create the tables.

In all the tables, the results that improve the non-distorted clustering error are marked with a double-box, and with a simple-box the results that maintain this clustering error. These boxes are included to focus the attention on the clustering error improvement.

Table 5.5 has many boxes because when the *MFW selection method* is applied, the best results are obtained. This is due to the fact that using this word selection method, the clustering is improved or maintained for every repository and every compression algorithm. These results are consistent with the ones shown in the clustering error figures, where it can be observed that the best clustering results correspond to the combination of the *MFW selection method* and the *asterisk substitution method*.

Fig 5.10 has been created to help and understand the tables. It compares Fig 5.5 and Table 5.5 with the aim of showing where the data in the table are from.

Table 5.5: *MFW selection method*. The clustering error obtained with no distortion, the minimum clustering error, and the maximum clustering error are shown. The frequencies when such clustering errors are obtained are shown as well. The results that improve the clustering error obtained with no distortion are highlighted inside a double-box. The results that maintain the non-distorted clustering error are highlighted inside a simple-box.

		Books		UCI-KDD		MedlinePlus		IMDB	
		Err	Freq	Err	Freq	Err	Freq	Err	Freq
LZMA	NoD	4	0.1-0.7	0	0.1-0.9	14	0.1-0.6	18	0.1-0.5
	Min	2	0.8,0.9	0	0.1-0.9	10	0.7-0.8	4	0.9
	Max	9	1	8	1	26	1	22	1
PPMZ	NoD	5	0.1-0.8	0	0.1-0.8	14	0.1-0.4,0.9	0	0.3-0.7,0.9
	Min	0	0.9	0	0.1-0.8	4	0.7	0	0.3-0.7,0.9
	Max	8	1	21	1	34	1	12	1
BZIP2	NoD	7	0.1-0.6	0	0.1-0.6	14	0.1-0.4,0.6	0	0.1-0.6
	Min	5	0.7	0	0.1-0.6	10	0.5,0.7-0.8	0	0.1-0.6
	Max	9	0.9	15	1	24	1	12	1

Table 5.6: *RW selection method*. The codification is the same as explained for Table 5.5.

		Books		UCI-KDD		MedlinePlus		IMDB	
		Err	Freq	Err	Freq	Err	Freq	Err	Freq
LZMA	NoD	4	0.1-0.6	0	0.1-0.6	14	-	18	-
	Min	4	0.1-0.6	0	0.1-0.6	13.8	0.3	13.3	0.7
	Max	9	1	8	1	28	1	22	1
PPMZ	NoD	5	0.1-0.2	0	0.1-0.3	14	-	0	-
	Min	5	0.1-0.2	0	0.1-0.3	14.2	0.3,0.5	5.2	0.2
	Max	10.7	0.9	21	1	34	1	12	1
BZIP2	NoD	7	-	0	-	14	-	0	-
	Min	5.9	0.8	1.6	0.2	14.6	0.3	8.4	0.1
	Max	10.5	0.1	17.2	1	24	1	15.3	0.4

Table 5.7: *LFW selection method*. The codification is the same as explained for Table 5.5.

		Books		UCI-KDD		MedlinePlus		IMDB	
		Err	Freq	Err	Freq	Err	Freq	Err	Freq
LZMA	NoD	4	-	0	0.1-0.2,0.5-1	14	-	18	0.1,0.5
	Min	9	0.1-0.4,0.6-1	0	0.1-0.2,0.5-1	20	0.1-0.2	10	0.6
	Max	12	0.5	8	1	28	0.5,0.7-1	22	0.2,0.4,0.7-1
PPMZ	NoD	5	-	0	-	14	-	0	-
	Min	7	0.1-0.2	15	0.3	20	0.1	8	0.1-0.3
	Max	11	0.3-0.6	21	0.5,0.7-1	34	0.7-1	12	0.6-1
BZIP2	NoD	7	-	0	-	14	-	0	-
	Min	4	0.3-0.4	8	0.3	16	0.1	10	0.6
	Max	9	0.1-0.2	16	0.1,0.6	26	0.3-0.4	32	0.1

Let us analyze Fig 5.10. The clustering error obtained when clustering the original documents, that is, the non-distorted clustering error, is 5. That is why there is a 5 in the cell that corresponds to the clustering error “Err” obtained with no distortion “NoD”. It can be observed that this clustering error remains constant from points 0.1 to 0.8. That is the reason why the cell that corresponds to the cumulative sum of frequencies “Freq” for the non-distorted results “NoD” is 0.1-0.8.

Similarly, the best clustering error obtained using the *asterisk substitution method* is 0, as can be seen looking at the point 0.9 in Fig 5.10. Therefore, the row that shows the minimum clustering error obtained - “Min” - shows a 0 in the cell that corresponds to the error “Err” and a 0.9 in the cell that corresponds to the cumulative sum of frequencies where this error is obtained “Freq”.




Figure 5.10: Understanding the tables.

Finally, the maximum clustering error obtained for the *asterisk substitution method* is 8. This error is obtained for a cumulative sum of frequencies of 1.0. Therefore, the table shows an 8 in the cell that corresponds to the maximum “Max” clustering error “Err” obtained, and it contains a 1 in the cell that corresponds to the cumulative sum of frequencies “Freq” for the maximum clustering error “Max”.

5.3.3 Synopsis of all the obtained results

Finally, this subsection summarizes all the obtained results in the form of four tables, one for each dataset. Table 5.8 corresponds to the Books dataset, Table 5.9 corresponds to the UCI-KDD dataset, Table 5.10 corresponds to the MedlinePlus dataset, and Table 5.11 corresponds to the IMDB dataset. The tables show the average clustering error for all the compression algorithms, all the word selection methods, and all the substitution methods. The clustering error is averaged as follows:

$$\text{Average CE} = \frac{\sum_{\forall \text{ distortion}} \text{CE}(\text{distortion})}{\# \text{ distortions}} \quad (5.1)$$

Where CE means “Clustering Error”, and the possible distortions go from 0.1 to 1.0. Therefore the number of distortions is always 10.

Analyzing these tables, one can reach the same conclusions as by analyzing the clustering error curves. However, summarizing the results by calculating the average clustering error helps to better see the differences between all the experiments carried out. Therefore, the tables shown in this subsection constitute an alternative and easier way of presenting the results obtained in the experiments developed in this chapter of the thesis.

Firstly, it can be observed that the average clustering error obtained using the *MFW selection method* is always less than the one obtained using the *RW selection method*, and this latter is always less than the one obtained using the *LFW selection method*.

Secondly, one can observe that in general, the average clustering error obtained applying the *asterisk substitution method* is less than the one obtained using the *random character substitution method*.

Finally, one can see that the best clustering error is obtained for a different compression algorithm depending on the dataset used. This could be due to the fact that each dataset is composed of texts of a different nature. The next chapter of the thesis tries to investigate the reasons why the non-distorted clustering error can be improved combining the *MFW selection method* and the *asterisk substitution method*.

Table 5.8: Books dataset. Average clustering error.

		MFW	RW	LFW
LZMA	Asterisks	4.10	5.51	9.30
	Random characters	14.80	23.42	28.00
PPMZ	Asterisks	4.80	7.44	9.00
	Random characters	15.10	21.83	23.90
BZIP2	Asterisks	7.90	7.99	7.00
	Random characters	21.10	26.39	31.70

Table 5.9: UCI-KDD dataset. Average clustering error.

		MFW	RW	LFW
LZMA	Asterisks	0.80	0.94	1.20
	Random characters	8.10	14.51	25.40
PPMZ	Asterisks	2.30	7.62	18.90
	Random characters	7.70	14.76	26.80
BZIP2	Asterisks	2.00	8.55	13.70
	Random characters	15.80	22.58	32.50

Table 5.10: MedlinePlus dataset. Average clustering error.

		MFW	RW	LFW
LZMA	Asterisks	14.40	16.98	25.40
	Random characters	16.40	19.28	26.80
PPMZ	Asterisks	13.80	18.58	28.60
	Random characters	16.20	19.62	27.60
BZIP2	Asterisks	14.20	17.88	23.00
	Random characters	19.00	24.20	30.20

Table 5.11: IMDB dataset. Average clustering error.

		MFW	RW	LFW
LZMA	Asterisks	13.40	16.24	19.40
	Random characters	20.70	23.79	31.90
PPMZ	Asterisks	2.60	8.18	10.40
	Random characters	10.70	19.51	25.10
BZIP2	Asterisks	2.60	12.05	17.40
	Random characters	20.50	25.09	33.50

5.4 Summary and Conclusions

This chapter of the thesis has taken a small step towards understanding both the nature of textual data and the nature of compression distances. This has been accomplished by performing an experimental evaluation of the impact that several kinds of word removal have on the NCD-based text clustering.

In terms of implementation, the CompLearn Toolkit [29], which implements the clustering algorithm described in [30, 75], has been used to carry out the experiments.

Six different distortion techniques have been evaluated. They are pairwise combinations of two factors: *word selection method* and *substitution method*. There are three word selection methods, depending on what words are chosen to be removed from the documents: *Most Frequent Words -MFW- selection method*, *Least Frequent Words -LFW- selection method* and *RW selection method*. There are two substitution methods, depending on the way in which the words are removed from the documents: *random character substitution method* and *asterisk substitution method*.

The NCD-driven clustering algorithm has been applied over four different datasets repeating the clustering three times using each time a different compression algorithm to calculate the NCD: PPMZ, LZMA and BZIP2.

In addition, in order to gain an insight into how the information is decreased when the distortion techniques are applied, the Kolmogorov complexity of the documents has been estimated based on the concept that data compression is an upper bound for it.

The experimental results have shown that the combination of the *selection method* and the *substitution method* is the key factor. Substituting the most frequent words using the *asterisk substitution method* is always the best option to maintain the most relevant information. In this case, the documents complexity estimation is slowly reduced and therefore the clustering error remains stable even though a considerable percentage of words were substituted from the documents. Moreover, its worth mentioning that, using the best distortion technique, even the non-distorted clustering error can be improved.

Analyzing Tables 5.5, 5.6, and 5.7 one can observe that the best results are obtained using the LZMA compression algorithm. This could be due to the fact that this compressor captures the contextual information because of its design. Section 4.2 explains the implementation details of all the compressors used in this thesis. Here a summarized description of the compression algorithms used in this thesis is given.

The LZMA algorithm codifies the symbols using as a dictionary part of the input stream previously seen. The method is based on a sliding window

that the encoder shifts as the strings of symbols are being encoded. The window is divided in two parts:

- *Search buffer*: Part of the input stream previously seen. This is the current dictionary.
- *Look-ahead buffer*: The text yet to be encoded.

It is important to point out that practical implementations of this method use really long *search buffers* of thousands of bytes long, and small *look-ahead buffers* of tens of bytes long [103].

Therefore, this compression algorithm takes the contextual information into account because it uses part of the input stream previously codified to codify the data that have yet to be codified.

The PPMZ is an adaptive statistical compression algorithm which is based on an encoder that maintains a statistical model of the text. It considers the N symbols preceding the symbol being processed. Therefore, this compressor takes the contextual information into account. However, the main difference, in this respect, between the PPMZ and the LZMA is that the former only considers about 10 symbols preceding the symbol being codified, whereas the latter considers thousands of symbols preceding it.

The BZIP2 is a block-sorting compressor that uses different techniques to compress the data. Some of these techniques transform the input by moving the symbols being encoded. In particular, the Burrows-Wheeler Transform, and the Move-To-Front transform behave that way. Therefore, this compression algorithm destroys the contextual information, since it shuffles the symbols in the compression process.

As a result of all the above, the next chapter of the thesis, which analyzes the relevance of the contextual information, only uses the LZMA to calculate the NCD. However, a deeper study of the effects that the loss or the maintenance of the contextual information have on the accuracy of the clustering results, using different compression algorithms, constitutes a future work.

Summarizing, three main contributions have been presented in this chapter. First, new text representations have been analyzed and studied with the aim of giving new insights for the evaluation of the NCD. Second, a technique which reduces the complexity of the texts while preserving most of the relevant information has been presented. Third, experimental evidence of how to fine-tune the representation of the documents, in order to obtain better NCD-driven clustering results, has been provided.

Chapter 6

Relevance of contextual information

The previous chapter experimentally evaluates the impact that several word removal techniques have on compression-based text clustering. It shows that the application of a specific distortion technique can improve the non-distorted clustering results. Since that technique implies, not only the removal of words, but also the maintenance of the previous text structure, exploring the relevance of both factors becomes necessary in order to better understand the results. This chapter explores precisely that.

The main contributions of this chapter of the thesis can be briefly summarized as follows:

- Experimental evaluation of the relevance that the contextual information has in compression-based text clustering, in a word removal scenario.
- New perspectives for the evaluation and explanation of the behavior of compression distances, in relation to contextual information.

The chapter is structured as follows. Section 6.1 describes the distortion techniques explored. Section 6.2 describes the experimental setup. Section 6.3 gathers and analyzes the obtained results. Finally, Section 6.4 summarizes the conclusions drawn from the experiments carried out in this chapter of the thesis.

6.1 Distortion Techniques

Four different distortion techniques are explored in this chapter of the thesis. One of them was analyzed in the previous chapter, and consists of incrementally removing the most frequent words in the English language, as described in depth in Section 5.1.

In order to maintain the length and the place of appearance of the removed words, instead of simply erasing them, their characters are replaced using asterisks. The marks that the words leave on the texts after the distortion is precisely what is called contextual information throughout the thesis. In this chapter, this technique is called *Original sorting* distortion technique because no random sorting is carried out after substituting the words with asterisks. The rest of the distortion techniques consist of first applying that distortion technique, and then randomly sorting different parts of the distorted texts. The description of the new distortion techniques is as follows:

- *Randomly sorting contextual information*: after replacing the words using asterisks, the strings of asterisks are randomly sorted. That is, the remaining words are maintained in their original places of appearance, while the removed words are not. It is important to note that each string of asterisks is treated as a whole. That is, if a word such as “hello” is replaced by “*****” these asterisks always remain together. This method is created in order to study whether the structure of the contextual information is relevant or not. Fig 6.1(c) represents a sample of text distorted using this technique.
- *Randomly sorting remaining words*: after replacing the words using asterisks, the remaining words are randomly sorted. That is, the contextual information is maintained, while the remaining words structure is not. This method is created to evaluate the importance of the structure of the remaining words. A visual representation of the effects of applying this technique can be seen in Fig 6.1(d).
- *Randomly sorting everything*: after replacing the words using asterisks, both the strings of asterisks and the remaining words are randomly sorted. It should be pointed out that, in this case, the strings of asterisks are randomly sorted as a whole too. This method is created as a control experiment. See Fig 6.1(e) for a visual representation of this technique’s effects.

The graphical differences among the four distortion techniques explored in this chapter can be seen in Fig 6.1. This figure clarifies the way in which each distortion technique modifies the texts.

```

.....
anemia is a condition
in which the body
does not have enough
healthy red blood
cells. red blood
cells provide oxygen
to body tissues.
.....

```

(a) Original text

```

.....
anemia ** * *****
** ***** *** ****
**** *** *****
healthy *** *****
***** ** *****
***** oxygen **
**** tissues
.....

```

(b) Original sorting

```

.....
anemia ***** **
***** ***** **
**** ***** ****
healthy ** ** **** *
** ** ***** oxygen
*****
tissues
.....

```

(c) Randomly sorting
contextual information

```

.....
oxygen ** * *****
** ***** *** ****
**** *** *****
anemia *** *****
***** ** *****
***** tissues **
**** healthy
.....

```

(d) Randomly sorting
remaining words

```

.....
***** ** **** * **
***** ***** ****
healthy ** ***** **
tissues ***** ****
***** ** ** anemia
**** *****
oxygen **
.....

```

(e) Randomly sorting
everything

Figure 6.1: Text distortion techniques.

6.2 Experimental Setup

The experiments have been carried out using the same compression-based clustering algorithm used in the first part of the thesis. The detailed description of this algorithm can be found in Section 5.2.1. In this part of the thesis, only the LZMA compression algorithm is used to perform the NCD-driven document clustering because the best results were obtained using it in the previous chapter.

6.2.1 Datasets

Five different datasets composed of texts written in English have been used in the experiments. Although the detailed description of the datasets can be found in Appendix B, a summarized description of them can be found here:

- **Books dataset:** Fourteen classical books from universal literature, to be clustered by author.
- **UCI-KDD dataset:** Sixteen messages from a newsgroup, to be clustered by topic.
- **MedlinePlus dataset:** Twelve documents from the MedlinePlus repository, to be clustered by topic.
- **IMDB dataset:** Fourteen plots of movies from the Internet Movie Data Base -IMDB- to be clustered by saga.
- **SRT-serial dataset:** Sixty-nine scripts of different serials which have been obtained from [93], to be clustered by serial.

6.3 Experimental Results

A figure is shown for every dataset. In each figure, the clustering error obtained applying the *Original sorting* distortion technique is plotted in the panel (a). In addition, in order to ease the comparison between this technique and the new distortion techniques, this curve is also plotted in the panels (b), (c) and (d), which correspond to the *Randomly sorting contextual information*, *Randomly sorting remaining words*, and *Randomly sorting everything* distortion techniques, respectively. Since these distortion techniques are based on randomly sorting different parts of the texts, the experiments have been repeated several times, and the mean and the standard deviation of the clustering error are depicted in panels (b), (c) and (d).

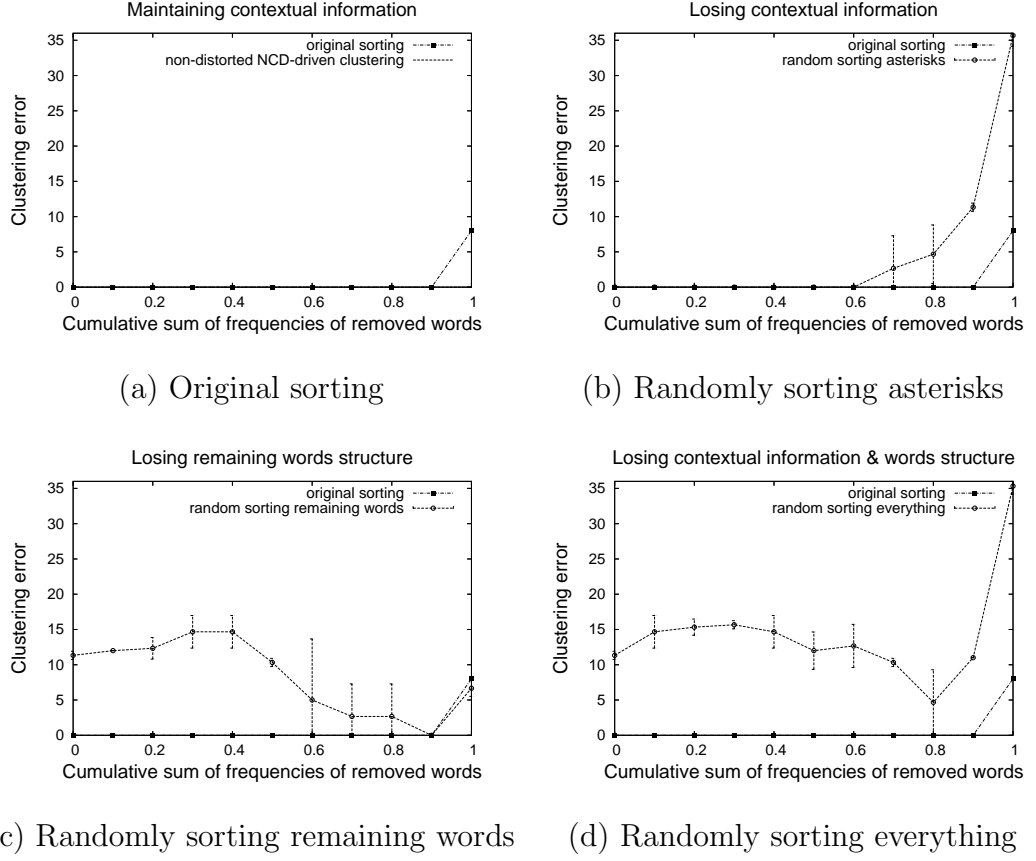


Figure 6.2: Clustering results for the UCI-KDD dataset.

In all the plots, the values on the vertical axis correspond to the obtained clustering error, while the values on the horizontal axis correspond to the cumulative sum of frequencies of the words that are removed from the texts.

Fig 6.2 shows the results that correspond to the UCI-KDD dataset. When the contextual information is not lost, the clustering error remains constant from 0.0 to 0.9, as can be observed looking at the panel (a). It is important to note that, in this case, the non-distorted clustering error cannot be improved since its value is 0.

Interesting conclusions can be drawn comparing that curve with the others. First, losing the contextual information makes the clustering results get worse as the amount of removed words increases. Second, losing the remaining words structure, the clustering results are worse when the texts contain many remaining words and few of contextual information. Third, losing every structure, both behaviors are observed at the same time, that is,

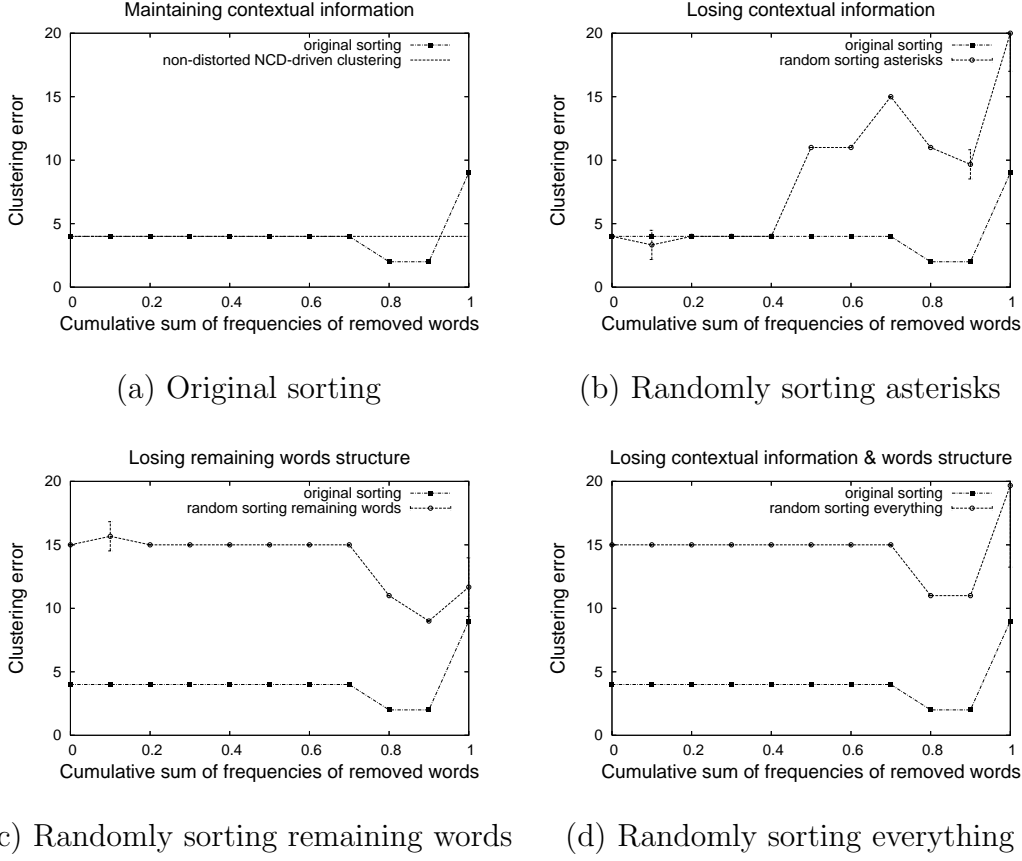


Figure 6.3: Clustering results for the Books dataset.

the clustering results are worse for small and big numbers of removed words. All these phenomena can be observed in panels (b), (c) and (d) respectively.

Fig 6.3 shows the results that correspond to the Books dataset. The curves show that the behavior of the distortion techniques is qualitatively similar to that observed for the UCI-KDD dataset. That is, when the contextual information is lost, the clustering error increases as the amount of removed words increases. When only the remaining words structure is lost, the clustering error is worse for small quantities of removed words. Look at panels (b) and (c) to observe this.

It is also important to mention that the non-distorted clustering error, depicted in (a) as a constant line, is only improved when the contextual information is maintained. Look at points 0.8 and 0.9 of the curve plotted in the panel (a) to notice this.

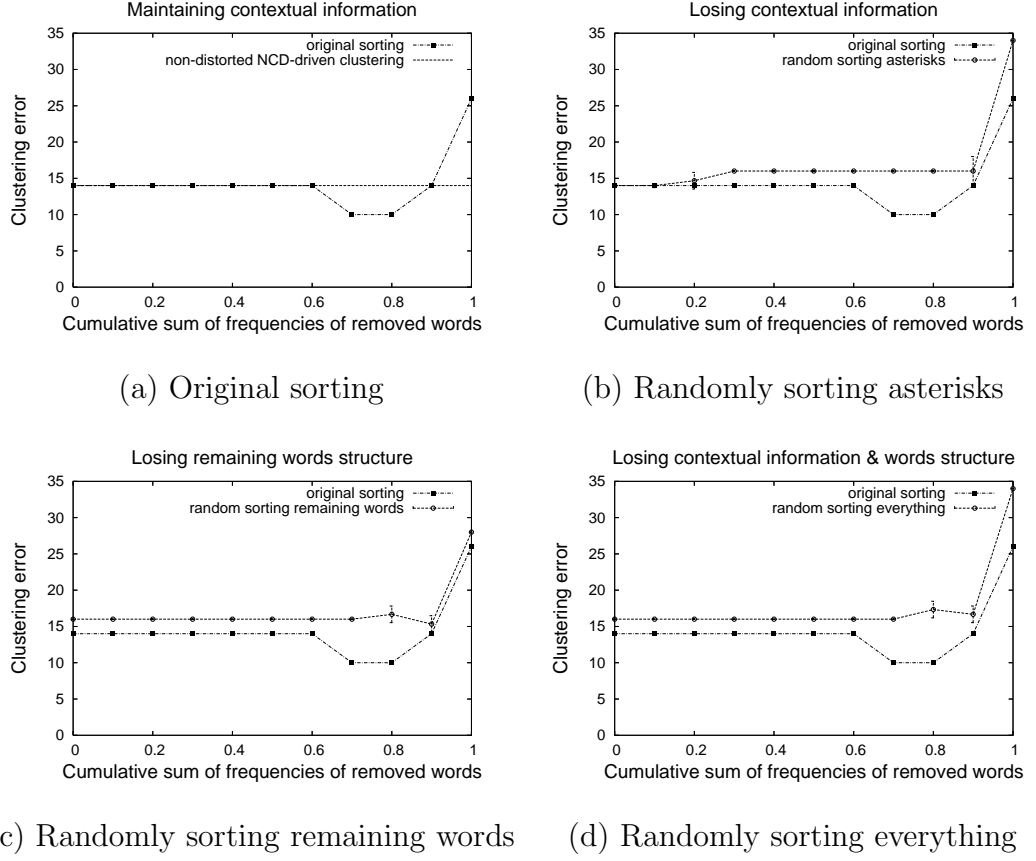


Figure 6.4: Clustering results for the MedlinePlus dataset.

Fig 6.4 depicts the results that correspond to the MedlinePlus dataset. The behavior shown in this figure is qualitatively similar to the previous one.

The results obtained for the IMDB dataset are depicted in Fig 6.5. The nature of this dataset is explained before analyzing the curves in order to better understand these interesting results. This dataset is composed of plots of movies to be clustered by saga. This means that as the amount of removed words increases, the words that still remain in the texts are words highly related to the sagas, such as for example names of characters or names of places. That is the reason why the non-distorted clustering error can be improved as much as panel (a) shows.

The clustering error in the cases in which the remaining words are randomly sorted are worse than the ones obtained when they are not. This means that the structure of the remaining words is highly relevant to this dataset. Nevertheless, when only the contextual information is lost, the clus-

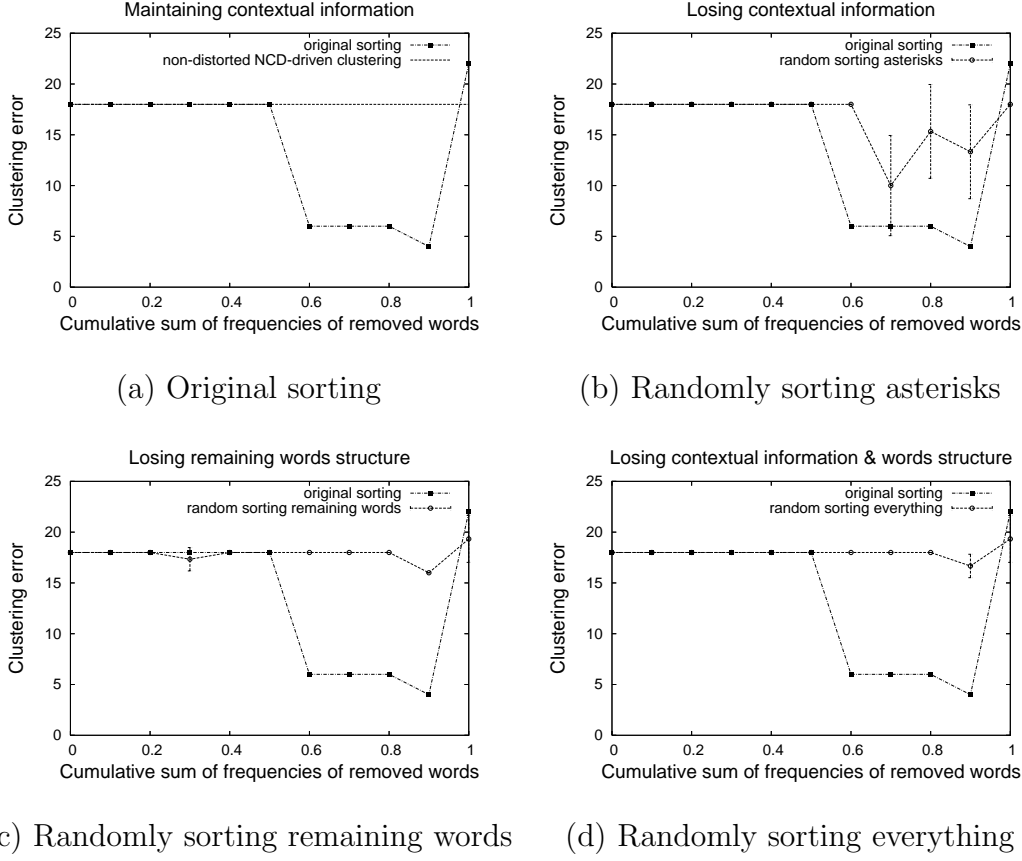


Figure 6.5: Clustering results for the IMDB dataset.

tering error is worse than when nothing is mixed up. This can be seen in panel (b). Therefore, it can be concluded that, although for this dataset the most relevant information corresponds to the remaining words, the contextual information is also relevant.

Finally, Fig 6.6 shows the results that correspond to the SRT-serial dataset. This is the bigger dataset that has been used in this chapter. Its results are consistent with the ones obtained for the rest of the datasets. That is, the contextual information is relevant as well, although the structure of the remaining words is more relevant than the contextual information when the amount of removed words, and therefore, the amount of contextual information is small.

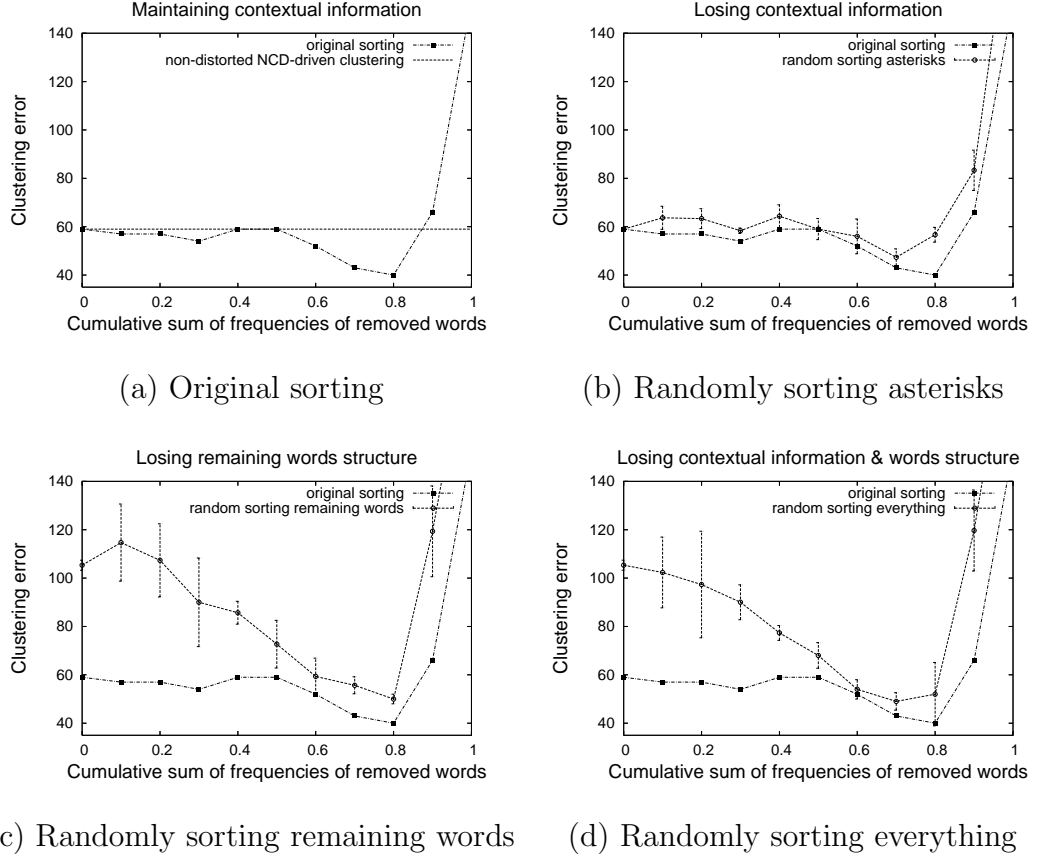


Figure 6.6: Clustering results for the SRT-serial dataset.

6.3.1 Synopsis of results

Additionally, Fig 6.7 has been created in order to better analyze the difference between the most interesting distortion techniques. It depicts the clustering error difference between the *Randomly sorting contextual information* and the *Randomly sorting remaining words* distortion techniques with respect to the *Original sorting* distortion technique. The length of each bar corresponds to the relative error, which is as follows

$$\Delta e_k = e_k - e_0, \quad (6.1)$$

where e_k is the clustering error obtained using the k distortion technique, e_0 is the clustering error obtained using the *Original sorting* distortion technique, and Δe_k is the relative error for the k distortion technique with respect to the *Original sorting* distortion technique.

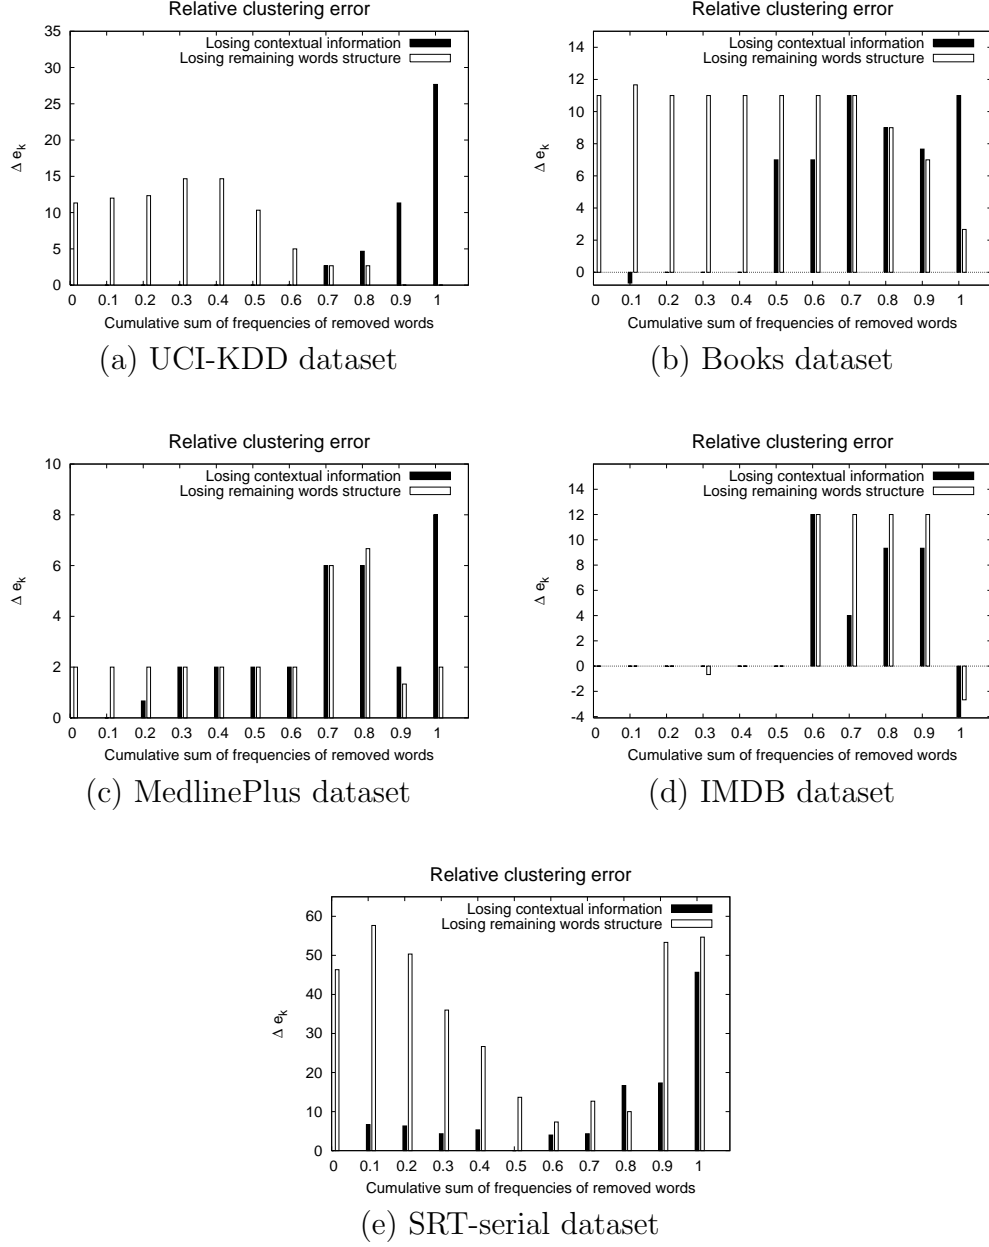


Figure 6.7: Clustering error difference with the *Original sorting* distortion technique. There exists a clustering error difference for both techniques. Therefore, one can conclude that both the remaining words structure and the contextual information are relevant in this scenario.

Fig 6.7 can be easily understood by analyzing the results obtained for the UCI-KDD dataset. These results can be seen in Fig 6.2 and in Fig 6.7(a). Looking at Fig 6.2(b)(c), two phenomena can be noticed. First, when the contextual information is lost, the clustering error gets worse from 0.7 to 1.0, see 6.2(b). Second, when the remaining words structure is lost, the clustering error is worse from 0 to 0.8, see 6.2(c). This is represented in Fig 6.7 in form of black and white bars: there are black bars from 0.7 to 1.0 and white bars from 0 to 0.8.

Analyzing the five plots depicted in Fig 6.7, it can be concluded that both the contextual information and the remaining words structure are relevant, since there exists a clustering error difference for both techniques with respect to the *Original sorting* method. In fact, Δe_k can be used to provide a quantitative measure of this relevance.

Finally, a summary of all the obtained results in the form of a table is shown. Table 6.1 contains the average clustering error for all the distortion techniques, and all the datasets used in this chapter of the thesis. The clustering error is averaged as follows:

$$\text{Average CE} = \frac{\sum_{\forall \text{ distortion}} \text{CE}(\text{distortion})}{\# \text{ distortions}} \quad (6.2)$$

Where CE means “Clustering Error”, and the possible distortions go from 0.1 to 1.0. Therefore the number of distortions is always 10.

Analyzing Table 6.1 one can reach the same conclusions as by analyzing the rest of the figures shown in the chapter. However, given that the table presents only a value for each pair *distortion technique-dataset*, comparing the effects that the distortion techniques have on the clustering error is easier analyzing the table than looking at the clustering error figures.

Comparing the first two rows of the table one can see that randomly sorting the contextual information increases the clustering error obtained when nothing is randomly sorted. In fact, if the average clustering error is normalized, then the difference between the distortion techniques can be more easily observed.

The clustering error can be normalized in the following manner:

$$E_k = \frac{e_k}{e_0}, \quad (6.3)$$

where E_k is the normalized average error obtained using the k distortion technique, e_0 is the average clustering error obtained using the *Original sorting* distortion technique, and e_k is the average clustering error obtained using the k distortion technique. A E_k greater than 1 implies an increment

of the average error. Therefore, it implies that the part of the texts which is randomly sorted by the k distortion technique is relevant for the NCD-driven text clustering.

Table 6.2 gathers the normalized average error for all the distortion techniques, with respect to the *Original sorting* distortion technique. Looking at the values presented in the table one can conclude that the contextual information is relevant for the NCD-driven clustering because when it is lost due to distortion, the average error gets worse, and therefore the normalized average error is greater than 1. This can be observed for all the datasets used in the chapter.

Table 6.1: Average clustering error. There is a difference in terms of average clustering error between the *Original sorting* distortion technique, and the rest of the distortion techniques. This means that both, the contextual information and the remaining words structure are relevant for the NCD-driven text clustering.

	UCI-KDD	Books	MedlinePlus	IMDB	SRT-serial
Original sorting	0.80	4.10	14.40	13.40	64.00
Randomly sorting contextual information	5.50	9.30	17.40	16.60	75.07
Randomly sorting remaining words	7.00	13.73	17.00	17.80	96.23
Randomly sorting everything	15.80	14.67	18.00	18.20	92.57

Table 6.2: Normalized average error. Analyzing these values one can reach the same conclusions as by analyzing the average error values. That is, both the contextual information and the remaining words structure are relevant for the NCD-driven text clustering.

	UCI-KDD	Books	MedlinePlus	IMDB	SRT-serial
Randomly sorting contextual information	6.88	2.27	1.21	1.24	1.17
Randomly sorting remaining words	8.75	3.35	1.18	1.33	1.50
Randomly sorting everything	19.75	3.58	1.25	1.36	1.45

6.4 Summary and Conclusions

The analysis that has been made in this chapter is the natural continuation of the one made in the previous chapter. In that chapter, different word removal techniques were applied to gradually filter the information contained in several sets of documents. It was shown that the application of a specific word removal technique could improve the non-distorted clustering results.

It is worth recalling that this technique was designed bearing the intrinsic nature of texts in mind. Texts contain words that provide a lot of information about the subject matter, at the same time as they contain other words with little meaning or relevance. Although, in principle, the non-relevant words are not as important as the relevant ones, the former constitute the substrate that supports the latter.

Generally, the more frequent a word is, the less relevant to the subject matter [79]. The main idea of the above mentioned distortion technique is to remove the non-relevant words maintaining the relevant ones. Thus, the distortion technique consists of removing the most frequent words in the English language from the documents. Instead of simply deleting the words, the technique replaces them using asterisks with the aim of maintaining the text's structure despite the removal. This simple idea allows maintenance of part of the contextual information, while filtering the information contained in the documents.

The immediate conclusion that can be drawn from the results that have been presented in Chapter 5 is that the words that remain in the documents after the application of such a distortion technique are the ones that contain the most relevant information in the texts. However, this distortion technique implies, not only the presence of some words, but also the presence of the previous text structure. Therefore, analyzing how the maintenance of the previous text structure affects the obtained results becomes necessary.

A comparison between this technique and three new distortion techniques that destroy the contextual information in different manners has been carried out in this chapter. Two main conclusions can be drawn by analyzing the results. First, it seems that maintaining the contextual information allows one to obtain better clustering results than losing it. Thus, it seems that by preserving the contextual information, the compressor is able to better capture the internal structure of the texts. Consequently, the compressor obtains more reliable similarities, and the non-distorted clustering results can be improved. Second, losing the structure of the remaining words affects the clustering results negatively. Therefore, it can be concluded that, in this scenario, both contextual information and remaining words have some relevance in the text clustering behavior.

Summarizing, two main contributions have been presented in this chapter of the thesis. First, the relevance that the contextual information has in this compression-based text clustering scenario has been evaluated. Second, new insights for the evaluation and explanation of the behavior of the compression distances, in relation to contextual information have been given.

Chapter 7

Application to Document Searching

The two previous chapters perform an experimental evaluation of the impact that several text distortion techniques have on the NCD-driven text clustering. They show that the application of a specific distortion technique can improve the non-distorted clustering results.

In this chapter, this distortion technique is applied to NCD-driven document search. It is worth mentioning that the document search method used in this chapter applies passage retrieval to address the problem that the NCD has when it is used to compare very different sized objects [30, 83].

The results presented in this chapter show that the application of the above mentioned distortion technique can improve the non-distorted search results.

The main contributions of this chapter of the thesis can be briefly summarized as follows:

- Practical application of the main conclusions taken from the studies developed in the first two parts of the thesis to document search.
- Improvement in the representation of documents that allows increasing the accuracy of the results obtained when searching documents.

The chapter is structured as follows. Section 7.1 describes the NCD-based document search method used in this part of the thesis. Section 7.2 describes the datasets used to perform the experiments. Section 7.3 gathers and analyzes the obtained results. Finally, Section 7.4 summarizes the conclusions drawn from the experiments presented in this chapter.

7.1 NCD-based Document Search Method

The NCD has been successfully applied to a wide range of domains, as stated in Chapter 4. Nevertheless, there is an issue that needs to be addressed if one wants to apply it under particular circumstances. Its drawback is that it does not commonly perform well when the compared objects are very different in size [30], as described in Section 4.3.

Although in many domains this does not constitute an obstacle, it can be a problem in those kinds of scenarios in which two objects of very different size are compared. This is, for example, the case of a typical document search scenario, because the size of the query can be very different from the size of the documents to be searched.

Because of that, in principle, the application of compression distances to document search cannot seem very appropriate. However, addressing this weakness, one can benefit from NCD's strengths. The NCD-based document search method used in this thesis [52, 82, 83] faces this problem by applying the philosophy of passage retrieval [105, 54, 20].

Passage retrieval is in principle similar to document retrieval, but involves the additional, preliminary stage of extracting passages from documents [62]. Thus, passage retrieval is based on considering the documents as sets of passages instead of considering them as atomic units.

Previous research has shown that passage retrieval can be used to improve document retrieval accuracy when the documents are long, have a complex structure, or are short but span many subjects [20].

There is a wealth of literature about different strategies which have been used to divide the documents into fragments. Thus, among others, structural features [105, 134, 145], or semantic features [54, 63, 86] have been used to delimit the passages. Another approach that consists of partitioning the documents into fragments of text of a given size has also been used [20, 62, 122, 135].

This last method, which is commonly called window-based, is used in this thesis because this approach is the most appropriate to solving the problem that the NCD has when the compared objects are very different in size. This is due to the fact that dividing the documents into fragments of equal size, simply avoids the problem.

Different sizes of windows have been used in window-based passage retrieval. For example, the use of passages of 150-300 words has been proposed in [20]. Other researchers have proposed passages exceeding 500 words [70]. More recent approaches have experimented with different lengths and a window size of 50 gave the best results in [135], while a bigger window size (200-1000) gave the best results in [122].

Since the search method used tries to solve the NCD weakness without restricting the size of the queries, it uses windows of different sizes. In particular, the sizes of windows go from 1 KB to N KB.

The minimum size of a window is 1KB due to the fact that the search engine is based on the NCD, and the NCD uses compression algorithms to estimate the entropy of the file, and compressors that are entropy encoders need the input file to be large in order to behave like an entropy encoder [103].

The maximum size N of a window depends on the compressor used. This is due to the fact that compression algorithms use a memory window which defines the best -most compressive- behavior of the algorithm. The engine uses a version of the Lempel-Ziv algorithm of a window size of 32 KB, therefore, $N = 32$. The engine stores the obtained passages in different databases, depending on their size. Thus, there are 32 different databases, since the documents are split in passages from 1 KB to 32 KB.

In the segmentation process, relevant paragraphs can be cut up and divided among different passages. This can lead to a critical fragmentation of the information contained in the paragraphs. The NCD-based document search method solves this problem by using overlap. Thus, each passage contains some bytes of the previous one. Further information on said NCD-based document search method can be found in [52, 82, 83].

7.2 Datasets

Three datasets composed of texts written in English have been used in the experiments. Although the detailed description of the data sets can be found in Appendix B, a summarized description of them can be found here:

- **UCM dataset:** 104 articles related to computer science written by researchers at the “Universidad Complutense de Madrid” -UCM- to be clustered by topic.
- **Reuters dataset:** 200 documents from a newsgroup from Reuters, to be clustered by topic. This dataset has been adapted to make it suitable for our experiments using the method described in Appendix B.
- **20newsgroups dataset:** This well-known dataset is composed of 20.000 documents on 20 different topics. The dataset can be downloaded from the UCI Knowledge Discovery in Databases Archive [125]. In the same way as the previous dataset, this dataset has been adapted

to make it suitable for our experiments using the method described in Appendix B.

Table 7.1 summarizes the characteristics of the dataset and the queries used in this chapter of the thesis. However, more detailed information on the datasets and the queries used can be seen in Appendices B and C, respectively.

Table 7.1: Datasets and experiments description.
(*) Adaptation described in Appendix B.

		UCM	<i>20newsgroups</i>	Reuters
Dataset description	#documents	104	20000	200
	#topics	11	20	10
Experiments description	documents	papers	one file per topic (*)	one file per topic (*)
	queries	abstracts	messages	news
	#queries	4	10	10
	queries' sizes	2 x 1KB 2 x 2KB	7 x 2KB 3 x 3KB	10 x 2KB

7.3 Experimental Results

The objective of this chapter is not to find the best way of retrieving information, but to show that some information can be more accurately retrieved by applying a distortion technique that changes the representation of the input data in a specific manner.

In a retrieval system, precision is the fraction of retrieved instances that are relevant. For example, if the system retrieves 10 results but only 6 of them are relevant, then the precision of the retrieved results would be 0.6.

Another measure commonly used in retrieval systems is the precision-at-K, which is the precision obtained for the first K retrieved results.

All the figures presented in this section contain a graph and a table. The graphs depict the precision-at-K obtained for the first K retrieved results, whereas the tables show the precision-at-K values for a selection of Ks. All the results are averaged over the queries used for each dataset.

There is one figure for each dataset. Each figure shows the benefits of applying distortion in a specific dataset. Fig 7.1 corresponds to the UCM dataset, Fig 7.2 corresponds to the *20newsgroups* dataset, and Fig 7.3 corresponds to the Reuters dataset.

The values on the “Distortion” axis correspond to the cumulative sum of the BNC-based frequencies of the words deleted from the documents. As

Section 5.1 explains, ten different sets of words to be removed from the documents are used to evaluate the distortion. Each set contains the words that account for a specific frequency, these values going from 0.1 to 1.0. The results obtained for each set correspond to the values from 0.1 to 1.0 in the “Distortion” axis. Moreover, the results obtained with no distortion are depicted in the figure as the values that correspond to a distortion of 0.

In addition, in order to ease the comparison between the precision values obtained, a summary of these values is shown in the tables contained in all the figures. Each row of a table contains the precisions obtained using a specific distortion. Thus, the first row contains the precision at K obtained without distortion, the second row corresponds to a distortion of 0.1, the third row corresponds to a distortion of 0.2, and so on. The precision values that maintain or improve the ones obtained with no distortion are highlighted in bold.


Looking at Fig 7.1 one can see that higher precision values are obtained when distorting the documents than when not distorting them for the UCM dataset. This fact can be noticed by comparing the non-distorted results and the rest of the results both in the surface and in the table. Therefore, one can assert that the application of document distortion improves the quality of the retrieved results for this dataset. In particular, the best precision value achieved without distortion is 0.70, while the best precision achieved using distortion is 0.80. This implies an improvement of 14%. Fig 7.4 summarizes the improvements obtained for the best distortions with respect to the results obtained without distortion.

Analyzing the results that correspond to the *20newsgroups* dataset -Fig 7.2- one can observe that, again, the results obtained when distorting the documents are better than the ones obtained when not distorting them. In particular, the best precision value achieved without distortion is 0.46, while the best precision achieved using distortion is 0.66. This implies an improvement of 43%.

The same happens in the Reuters dataset. Thus, looking at Fig 7.3 one can see that the best precision value obtained without distortion is 0.42, whereas the best precision obtained using distortion is 0.46. Although this improvement is lower than the one achieved for the *20newsgroups* dataset, it implies an improvement of 10%.


7.3.1 Summary of Results

The figures presented in the previous subsection comprehensively show all the experimental results. That is, they show the precision values obtained for all the explored distortions. In this subsection, a figure that summarizes




Dist.	Average precision at K passages							
	5	10	15	20	30	40	50	100
0.0	0.70	0.60	0.62	0.53	0.49	0.43	0.41	0.32
0.1	0.80	0.80	0.75	0.73	0.63	0.58	0.58	0.46
0.2	0.80	0.73	0.72	0.66	0.64	0.63	0.57	0.47
0.3	0.75	0.78	0.73	0.71	0.63	0.58	0.56	0.47
0.4	0.80	0.80	0.72	0.68	0.65	0.59	0.58	0.47
0.5	0.80	0.75	0.73	0.74	0.69	0.66	0.63	0.52
0.6	0.80	0.80	0.80	0.74	0.64	0.61	0.57	0.50
0.7	0.80	0.80	0.77	0.75	0.64	0.63	0.60	0.49
0.8	0.70	0.55	0.48	0.49	0.49	0.49	0.47	0.37
0.9	0.70	0.63	0.50	0.48	0.42	0.36	0.33	0.29
1.0	0.45	0.30	0.27	0.28	0.27	0.23	0.22	0.21

Figure 7.1: UCM dataset. Benefits of applying distortion. The precision values that maintain or improve the ones obtained with no distortion are highlighted in bold.



Dist.	Average precision at K passages							
	5	10	15	20	30	40	50	100
0.0	0.46	0.33	0.31	0.27	0.23	0.20	0.19	0.15
0.1	0.60	0.47	0.42	0.38	0.32	0.29	0.25	0.27
0.2	0.62	0.47	0.43	0.40	0.32	0.27	0.24	0.19
0.3	0.64	0.47	0.42	0.37	0.32	0.28	0.25	0.20
0.4	0.66	0.52	0.46	0.40	0.33	0.28	0.26	0.19
0.5	0.64	0.51	0.42	0.37	0.31	0.28	0.25	0.20
0.6	0.60	0.48	0.43	0.40	0.33	0.28	0.26	0.19
0.7	0.44	0.37	0.37	0.34	0.27	0.23	0.21	0.16
0.8	0.42	0.35	0.29	0.25	0.20	0.19	0.17	0.13
0.9	0.24	0.26	0.20	0.17	0.15	0.16	0.15	0.11
1.0	0.02	0.12	0.11	0.09	0.10	0.09	0.09	0.08

Figure 7.2: *20newsgroups* dataset. Benefits of applying distortion. Again, the precision values that maintain or improve the ones obtained with no distortion are highlighted in bold.



Dist.	Average precision at K passages							
	5	10	15	20	30	40	50	100
0.0	0.42	0.33	0.31	0.29	0.26	0.23	0.21	0.17
0.1	0.46	0.37	0.33	0.31	0.29	0.27	0.25	0.30
0.2	0.40	0.34	0.31	0.27	0.27	0.26	0.24	0.18
0.3	0.40	0.33	0.30	0.29	0.26	0.24	0.22	0.18
0.4	0.38	0.37	0.31	0.29	0.26	0.24	0.24	0.19
0.5	0.42	0.37	0.33	0.31	0.26	0.24	0.23	0.18
0.6	0.40	0.35	0.31	0.28	0.27	0.25	0.23	0.18
0.7	0.30	0.29	0.26	0.24	0.22	0.21	0.20	0.17
0.8	0.24	0.28	0.30	0.28	0.22	0.21	0.20	0.17
0.9	0.22	0.29	0.25	0.24	0.19	0.19	0.17	0.14
1.0	0.12	0.09	0.11	0.14	0.13	0.12	0.12	0.12

Figure 7.3: Reuters dataset. Benefits of applying distortion. The precision values that maintain or improve the ones obtained with no distortion are highlighted in bold.

those results is presented. This figure -Fig 7.4- depicts the percentage of improvement for the best distortion, with respect to the results obtained with no distortion.




Figure 7.4: Percentage of improvement for the best distortion with respect to the results obtained with no distortion.

Although the best precision is achieved at a different distortion depending on the dataset used, this is consistent with the cut-offs defined by Luhn which exclude non-significant words [79]. These cut-offs have to be established by trial and error because they are different for each dataset [126]. That is, the non-relevant words, and therefore, the relevant ones, depend on the dataset used.

However, analyzing the curves plotted in Fig 7.4 one can observe that although the percentage of improvement is different for each dataset, there is always an improvement. Therefore, it can be concluded that applying document distortion is advantageous for the search method used in this thesis in terms of accuracy.

7.4 Summary and Conclusions

This chapter of the thesis has applied the document distortion technique that was found to be beneficial in document clustering -see Chapters 5 and 6- to document search. In particular, an NCD-driven document search method based on passage retrieval has been used in this chapter.

The experimental results have shown that the non-distorted search results can be improved by applying such a document distortion technique. The results are consistent across all the datasets used, that is, the search accuracy has been improved in all cases, as Fig 7.4 evidences. However, the best precision achieved for each dataset has been obtained at a different distortion. This behavior is consistent with the cut-offs defined by Luhn, an upper and a lower, which exclude the non-significant words [79]. These cut-offs have to be established by trial and error [126], which means that they are different for each dataset. That is, the non-relevant words, and therefore, the relevant ones, depend on the dataset used. This can explain why the best precision is achieved at a different distortion for each dataset.

Summarizing, two main contributions have been presented in this chapter of the thesis. First, a practical application of the main conclusions taken from the studies developed in the first two parts of the thesis to a document search scenario has been given. Second, a representation of the documents that improves the non-distorted document search accuracy has been proposed.

Chapter 8

Conclusions

The first part of this thesis -Chapter 5- has taken a step towards understanding compression distances by performing an experimental evaluation of the impact that several kinds of word removal have on compression-based text clustering. Six different distortion techniques based on word removal have been applied to gradually filter the information contained in the datasets. See Section 5.1 for a detailed description of the distortion techniques.

It has been shown that the application of a specific word removal technique can improve the non-distorted clustering results. This fact can be observed analyzing the figures shown in Section 5.3 .

The distortion technique that leads to an improvement in the non-distorted clustering results consists in erasing the most frequent words in the English language from the documents, replacing their characters with asterisks. This strategy allows preservation of the text structure despite the word removal. The experimental results presented in Chapter 5 show that the application of this distortion technique improves the non-distorted clustering results even, and precisely when, a lot of words are removed from the documents.

Since this technique implies, not only the removal of words, but also the maintenance of the previous text structure, analyzing the impact of both factors becomes necessary in order to better understand why the clustering results can be improved by applying the technique. This research is carried out in the second part of the thesis, which corresponds to Chapter 6.

Thus, the analysis made in Chapter 6 is the natural continuation of the one made in Chapter 5. Chapter 6 explores how the loss or the maintenance of the contextual information affects the clustering accuracy. Moreover, it explores how the loss or the preservation of the remaining words structure affects the clustering.

This has been accomplished by creating three new distortion techniques based on the distortion technique, presented in Chapter 5, which improves

the clustering accuracy. The detailed description of said distortion techniques is shown in Section 6.1.

Two main conclusions can be drawn by analyzing the results presented in Section 6.3. Firstly, it seems that maintaining the contextual information allows one to obtain better clustering results than losing it. Thus, it seems that by preserving the contextual information, the compressor is better able to capture the internal structure of the texts. Consequently, the compressor obtains more reliable similarities, and the non-distorted clustering results can be improved. Secondly, losing the structure of the remaining words affects the clustering results negatively. Therefore, it can be concluded that, in this scenario, both contextual information and remaining words have some relevance in the text clustering behavior.

Everything learned in the first two parts of the thesis has been applied to compression-based document search in the last part of the thesis -Chapter 7-. The objective of that chapter is not finding the best way of retrieving information, but showing that some information can be more accurately retrieved by applying a distortion technique that changes the representation of the input data in a specific manner.

Despite the wide and successful use of compression distances, their application to document search is not trivial due to their having a weakness that must be taken into account if one wants to apply them to document search. Their drawback is that they do not commonly perform well when the compared objects have very different sizes. An NCD-based document search engine that deals with that drawback by using passage retrieval, is used in the last part of the thesis.

The results presented in Section 7.3 show that the non-distorted search results can be improved by applying the document distortion technique that improves the non-distorted clustering results. This fact augments the applicability of the distortion technique, presented in Chapter 5, which fine-tunes the text representation.

Summarizing, from all the results presented in this thesis, it can be concluded that the application of one of the explored distortion techniques can be beneficial both in NCD-based document clustering and in NCD-based document search.

Chapter 9

Summary of Results

This chapter shows how the main objectives of this thesis have been achieved.

- *The attainment of objective 1: Providing new perspectives for understanding the nature of textual data.*

All the experiments carried out in this thesis have been focused on better understanding of both the nature of textual data, and the nature of compression distances. The way in which the thesis studies textual data is evaluating the impact that diverse distortion techniques have on compression-based document clustering -Chapters 5 and 6-, and their impact on compression-based document search -Chapter 7-.

The experimental results presented in those chapters have shown that distorting the documents, by erasing the most frequent words in the English language in a specific manner, improves the accuracy of both the document clustering and the document search. This means that such distortion preserves the relevant information contained in the documents while removing the non-relevant information.

Moreover, the results presented in Chapter 6 have shown that the maintenance of the previous structure of the texts, despite the word removal has beneficial effects on the clustering results. Therefore, it seems that the contextual information is relevant in this kind of scenario.

- *The attainment of objective 2. Providing a technique to smoothly reduce the complexity of the documents while preserving most of their relevant information.*

As said previously, several distortion techniques have been evaluated in this thesis. One of them produces both a smooth decrease of the non-relevant information in the set of documents considered, and a smooth

decrease of the documents complexity estimation. The application of this technique leads to an improvement in both the document clustering accuracy and the document search accuracy. The detailed description of the distortion technique can be seen in Chapter 5. The effects of applying that technique can be seen in Sections 5.3, 6.3, and 7.3.

- *The attainment of objective 3. Giving experimental evidence of how to fine-tune the text representation so that better results are obtained when using the NCD-driven text clustering.*

The above mentioned distortion technique fine-tunes the text representation by filtering the information contained in the texts. This leads to better clustering results, as Sections 5.3, and 6.3 show.

- *The attainment of objective 4. Giving new insights for the evaluation and explanation of the behavior of the NCD.*

The distortion techniques explored in this thesis are the tool that allows one to study the behavior of the NCD. Chapters 5, 6, and 7 present such a study.

It seems that distorting the documents by removing the most frequent words in the English language, the compressor is able to better capture the internal structure of the texts. Consequently, the compressor obtains more reliable similarities, and the results can be improved.

- *The attainment of objective 5. Experimentally evaluating the relevance that the contextual information has in compression-based text clustering, in a word removal scenario.*

The distortion technique that fine-tunes the text representation implies not only the removal of words, but also the maintenance of the previous text structure. Chapter 6 has explored the relevance of both factors with the aim of better understanding the results. It has been observed that maintaining the contextual information allows one to obtain better clustering results than losing it, as Section 6.3 shows.

- *The attainment of objective 6. Applying the main conclusions taken from the studies developed in the first two parts of the thesis to document search.*

The technique that filters the information contained in the texts by removing the irrelevant words while maintaining the previous text structure can be used in different application domains. Chapter 7 applies it to document search with the purpose of studying if the improvements

observed when clustering documents are observed when searching documents, as well.

- *The attainment of objective 7. Giving a representation of documents that improves the non-distorted document search accuracy.*

Since text representation plays an important role in document search, the application of a distortion technique that fine-tunes the representation of texts can be beneficial. Chapter 7 has explored if the above mentioned distortion technique can lead to better document search results. The experimental results have shown that changing the representation of the documents by applying such a distortion technique leads to better results, as Section 7.3 shows.

Chapter 10

Contributions

- INTERNATIONAL JOURNALS

- **Reducing the Loss of Information through Annealing Text Distortion**, Ana Granados, Manuel Cebrián, David Camacho, and Francisco de Borja Rodríguez, *IEEE Transactions on Knowledge and Data Engineering*, vol.23, n. 7, July 2011. ISSN 1041-4347. JCR (2009): impact factor = 2.285; Ranking (Comput. Science-Inform. Systems) = 20/116.

This work explores different text distortion techniques based on word removal. It analyzes how the information contained in the documents and how the upper bound estimation of their Kolmogorov complexity progress as the words are removed from the documents in different manners. Three different ways of choosing the words to be removed and two different ways of removing the words are explored. Combining these two factors, six different distortion techniques are obtained, and therefore, six methods are explored in the work. A compression-based clustering method is used to experimentally evaluate the impact that the studied distortion techniques have on the amount of information contained in the distorted documents. The experimental results show that the application of a specific distortion technique can improve the clustering accuracy.

- **Is the contextual information relevant in text clustering by compression?**, Ana Granados, David Camacho, Francisco de Borja Rodríguez. *Submitted to Expert Systems with Applications in February 2011.*

This work is the natural extension of the previous one. While the previous work evaluates the impact that several kinds of word

removal have on compression-based text clustering, showing that the application of one of the explored distortion techniques can improve the non-distorted clustering results, this work analyzes the reasons why applying that technique can improve the clustering results. The technique implies not only the removal of words, but also the maintenance of the previous text structure. Therefore, exploring the relevance of both factors becomes necessary in order to better understand the results. That is precisely what this work analyzes. The experimental results show that maintaining the contextual information allows one to obtain better clustering results than losing it.

- **Improving the Accuracy of the Normalized Compression Distance combining Document Segmentation and Document Distortion**, , Ana Granados, Rafael Martínez, David Camacho, Francisco de Borja Rodríguez. *Submitted to IEEE Transactions on Knowledge and Data Engineering in October 2011*. This work applies the knowledge learned in the previous ones to the retrieval of documents. The retrieval approach is based on using document segmentation and document distortion. The experimental results show that the application of the previously explored distortion technique can improve the retrieval results.

- INTERNATIONAL CONFERENCES

- **Evaluating the Impact of Information Distortion on Normalized Compression Distance-driven Text Clustering**, Ana Granados, Manuel Cebrián, David Camacho, and Francisco de Borja Rodríguez, *In proceedings of 2nd International Castle Meeting on Coding Theory and Applications, ICMCTA 2008, Lecture Notes in Computer Science, Vol. 5228*

This work is the prelude to the work deeply developed in the journal *Reducing the Loss of Information through Annealing Text Distortion*. While the former only uses a compression algorithm and a dataset, the latter uses three compression algorithms, each of them belonging to a different family of compressors, and four different datasets.

- **Relevance of contextual information in compression-based text clustering**, Ana Granados, Rafael Martínez, David Camacho, Francisco de Borja Rodríguez, *In proceedings of 11th International Conference on Intelligent Data Engineering and Automated Learning, IDEAL 2010, Lecture Notes in Computer Science*

This work is the prelude to the work deeply developed in the journal *Is the contextual information relevant in text clustering by compression?*. While the former only uses two distortion techniques and two datasets, the last uses four distortion techniques, and five different datasets.

- **Influence of music representation on compression-based clustering**, Antonio González-Pardo, Ana Granados, David Camacho, Francisco de Borja Rodríguez, *In proceedings of the IEEE World Congress on Evolutionary Computation, IEEE CEC 2010*

This paper constitutes a parallel work to the main work developed in the thesis. The paper applies the Normalized Compression Distance to music clustering. The paper analyzes how the selection of a particular representation of music audio files can affect the clustering process. Three different music representations are explored in the paper: binary code, wave information, and SAX. In addition, two different algorithms are applied to automatically perform the clustering: a hierarchical clustering method based on the quartet tree method, and a genetic algorithm. The experimental results show how the representation of the music file plays a decisive role in the NCD-driven clustering. Thus, the best clustering results are obtained when the music is represented using its wave information. The other representations - WAV file, and SAX - get worse results. This can be due to the fact that the wave representation implies a smaller information loss. The paper also presents a new clustering method based on genetic algorithms as an alternative to the clustering method developed in the CompLearn Toolkit.

Appendix A

Acronyms

BNC British National Corpus

BWT Burrows-Wheeler Transform

BZIP Basic ZIPper

HMM Hidden Markov Model

IMDB Internet Movie Data Base

IT Information Theory

KNN K-Nearest Neighbour

LFW Least Frequent Words

LZ Lempel-Ziv

LZMA Lempel-Ziv-Markov chain Algorithm

MFW Most Frequent Words

MTF Move to Front Transform

NCD Normalized Compression Distance

NID Normalized Information Distance

PPM Prediction with Partial string Matching

RLE Run Length Encoding

RW Random Words

UCI-KDD University of California, Irvine, Knowledge Discovery in Databases archive

UCM Universidad Complutense de Madrid

USM Universal Similarity Metric

VSM Vector Space Model

Appendix B

Datasets

The detailed description of the datasets used throughout the thesis can be found here. All of the datasets comprise several texts written in English. An extract of one of the documents can be seen in a picture for every dataset.

B.1 Books dataset

Fourteen classical books, to be clustered by author. There are:

- Two books by Agatha Christie:
 - *The Secret Adversary*
 - *The Mysterious Affair at Styles*
- Three books by Alexander Pope:
 - *An Essay on Criticism*
 - *An Essay on Man*
 - *The Rape of the Lock, an heroic-comical Poem*
- Two books by Edgar Allan Poe:
 - *The Fall of the House of Usher*
 - *The Raven*
- Two books by Miguel de Cervantes:
 - *Don Quixote*
 - *The Exemplary Novels*

In a village of La Mancha, the name of which I have no desire to call to mind, there lived not long since one of those gentlemen that keep a lance in the lance-rack, an old buckler, a lean hack, and a greyhound for coursing. An olla of rather more beef than mutton, a salad on most nights, scraps on Saturdays, lentils on Fridays, and a pigeon or so extra on Sundays, made away with three-quarters of his income. The rest of it went in a doublet of fine cloth and velvet breeches and shoes to match for holidays, while on week-days he made a brave figure in his best homespun. He had in his house a housekeeper past forty, a niece under twenty, and a lad for the field and market-place, who used to saddle the hack as well as handle the bill-hook. The age of this gentleman of ours was bordering on fifty; he was of a hardy habit, spare, gaunt-featured, a very early riser and a great sportsman. They will have it his surname was Quixada or Quesada (for here there is some difference of opinion among the authors who write on the subject), although from reasonable conjectures it seems plain that he was called Quexana. This, however, is of but little importance to our tale; it will be enough not to stray a hair's breadth from the truth in the telling of it.

Figure B.1: Books. Extract from *Don Quixote* by Miguel de Cervantes.

- Three books by Niccolò Machiavelli:
 - *Discourses on the First Decade of Titus Livius*
 - *History of Florence and of the Affairs of Italy*
 - *The Prince*
- Two books by William Shakespeare:
 - *The tragedy of Antony and Cleopatra*
 - *Hamlet*

Since the documents belonging to this dataset are books, their size is quite big in general.

B.2 UCI-KDD dataset

Sixteen messages from a newsgroup (UCI-KDD) [125], to be clustered by topic. There are:

- Three documents on atheism.

In case people think email scanning doesn't take place, I can assure you that it is done regularly by many sites - usually not by government agencies (or at least not that I know of), but by local administrators who, for reasons of their own, have decided to monitor all communications (I'm sure you can all think of a whole mess of reasons - stop hackers/ terrorists/child pornographers/drug dealers/evil commies/whatever). There have been several occasions when I've got people into trouble simply by including the traditional NSA bait in a message (I don't try it any more now :-). A friend of mine was once picked up for mentioning the name of the UK town of Scunthorpe (hint: look for words embedded in it). I'm sure there are many more examples of this happening (in fact if anyone has any examples I'd appreciate hearing from them - I could use them as ammunition during talks on privacy issues).

Figure B.2: UCI-KDD. Extract from a document on cryptography.

- Three documents on Christian religion and homosexuality.
- Two documents on Christian religion and reincarnation.
- Two documents on politics and guns.
- Three documents on cryptography, governments and communications.

- Three documents on inherent problems of cryptography.

The main characteristic of these texts is their small size.

B.3 MedlinePlus dataset

Twelve documents from the MedlinePlus repository [84], to be clustered by topic. There are:

- Three documents related to alcohol:
 - Alcohol use
 - Alcoholic neuropathy
 - Alcoholism
- Three documents on diabetes:
 - Diabetes diet
 - Diabetes education

In many cases, moderate weight loss and increased physical activity can control type 2 diabetes. Some people will need to take oral medications or insulin in addition to lifestyle changes.

Children with type 2 diabetes present special challenges. Meal plans should be recalculated often to account for the child's change in calorie requirements due to growth. Three smaller meals and 3 snacks are often required to meet calorie needs.

Changes in eating habits and increased physical activity help reduce insulin resistance and improve blood sugar control. When at parties or during holidays, your child may still eat sugar-containing foods, but have fewer carbohydrates on that day. For example, if birthday cake, Halloween candy, or other sweets are eaten, the usual daily amount of potatoes, pasta, or rice should be eliminated. This substitution helps keep calories and carbohydrates in better balance.

For children with either type of diabetes, special occasions (like birthdays or Halloween) require additional planning because of the extra sweets.

Figure B.3: MedlinePlus. Extract from a document on diabetes.

- Diabetes definition
- Three documents on meningitis:
 - Meningitis gramnegative
 - Meningitis meningococcal
 - Meningitis staphylococcal
- Three documents on tumors:
 - Hepatocellular carcinoma
 - Spinal tumor
 - Thyroid cancer

Since these texts are about medicine, they are very specific and their vocabulary is very specialized.

B.4 IMDB dataset

Fourteen plots of movies from the Internet Movie Data Base (IMDB) [60], to be clustered by saga. There are five different sagas:

- Indiana Jones:
 - *Raiders Of The Lost Ark*
 - *Temple Of The Doom*
 - *The Last Crusade*
- Pirates Of The Caribbean
 - *The Curse of the Black Pearl*
 - *Dead Man's Chest*
 - *At World's End*
- The initial saga of Star Wars
 - *A New Hope*
 - *The Empire Strikes Back*
 - *Revenge of the Jedi*
- The Matrix
 - *The Matrix*
 - *Matrix Reloaded*
 - *Matrix Revolutions*
- The Mummy
 - *The Mummy*
 - *The Mummy Returns*

An important characteristic of these documents is the presence of names of characters and places that are related to the sagas.

B.5 SRT-serial dataset

Sixty-nine scripts of different serials which have been obtained from [93], to be clustered by serial. There are three chapters of each of these serials:

- *Accidentally on purpose*
- *Bones*
- *Community*

Thomas A. Anderson is a man living two lives. By day he is an average computer programmer and by night a malevolent hacker known as Neo. Neo has always questioned his reality but the truth is far beyond his imagination. Neo finds himself targeted by the police when he is contacted by Morpheus, a legendary computer hacker branded a terrorist by the government. Morpheus awakens Neo to the real world, a ravaged wasteland where most of humanity have been captured by a race of machines which live off of their body heat and imprison their minds within an artificial reality known as the Matrix. As a rebel against the machines, Neo must return to the Matrix and confront the agents, super powerful computer programs devoted to snuffing out Neo and the entire human rebellion.

Figure B.4: IMDB. Extract from the movie *The Matrix*.

- *CSI New York*
- *Damages*
- *Dexter*
- *Doctor Who*
- *Eastwick*
- *Emergency room*
- *Heroes*
- *House*
- *How I met your mother*
- *Justified*
- *Law and order*
- *Lost*
- *New tricks*
- *Northern exposure*
- *Nurse Jackie*
- *Parenthood*
- *Supernatural*
- *The life and times of Tim*
- *Till death*
- *Ugly Americans*

My name is Chandra Suresh. I'm a geneticist.
I have a theory about human evolution, and I believe you are a part of it.
What makes some walk a path of darkness, while others choose the light?
Is it will?
Is it destiny?
Can we ever hope to understand the force that shapes the soul?
For thousands of years, my people have taken spirit walks, following destiny's path into the realms of the unconsciousness.
I'm ready to begin my journey.
To fight evil, one must know evil.
One must journey back through time and find that fork in the road.
Where heroes turn one way, and villains turn another.

Figure B.5: SRT-serial. Extract from the serial *Heroes*.

The nature of this dataset is similar to the previous one because the documents contain names of characters and places that are related to the serials.

B.6 UCM dataset

This dataset is composed of 104 articles related to computer science written by researchers at the *Universidad Complutense de Madrid* (UCM). All the articles have been extracted from the UCM Computer Science Department website (<http://www.fdi.ucm.es/investigacion>). Then, they have been carefully classified in different knowledge areas:

- Grid computing
- Architecture and Technology of Computing Systems
- Cloud computing
- Finance application
- Software Engineering applied to e-Learning
- Software Agents
- Semantic Web
- Declarative Programming
- Natural Language Processing
- Petri Nets

The financial services industry today produces and consumes huge amounts of data and the processes involved in analysing these data have large and complex resource requirements. The need to analyse the data using such processes and get meaningful results in time, can be met only up to a certain extent by current computer systems. Most service providers attempt to increase efficiency and quality of their service offerings by stacking up more hardware and employing better algorithms for data processing. However, there is a limit to the gains achieved by using such an approach. One viable alternative would be to use emerging technologies such as the Grid. Grid computing and its application to various domains have been actively studied by many groups for more than a decade now. In this paper we explore the use of the Grid in the financial services domain; an area which we believe has not been adequately looked into.

Figure B.6: UCM. Extract from a paper on grid computing.

- Real-Time Systems

In this case, the documents stored in the databases correspond to articles, while the queries used in the experiments consist of abstracts of other articles related to the above ones.

It is important to note that the size of the documents (articles) and the queries (abstracts) is very different. This fact affects the behavior of the NCD because the NCD does not fit well when the compared objects are very different in size. This kind of framework is very suitable for evaluating the NCD-based document retrieval method used in Chapter 7.

B.7 Reuters dataset

This dataset is composed of 200 documents from the well-known Reuters-21578 corpus, which contains texts on 10 different topics. The whole dataset can be downloaded from kdd.ics.uci.edu/databases/reuters21578/reuters21578.html.

It should be pointed out that most of the documents contained in this dataset have a similar size. Therefore, in principle, this dataset does not seem very suitable for evaluating the NCD-based document retrieval method used in the third part of this thesis. However, it has been adapted to make it suitable for the experiments. The adaptation carried out consists of creating one big file per topic in the following manner:

- First, the documents that will constitute the queries are randomly selected among all the documents contained in the dataset.
- Then, the rest of the documents are used to create one big file per topic by concatenating all the documents concerning a topic. Note that the documents selected to build the queries are not included in that big file.

This adaptation makes the size of the documents and the size of the queries very different. In this way, the dataset becomes suitable for evaluating the NCD-based document retrieval method used.

- Acquisitions
- Corn
- Crude
- Earn
- Grain
- Interest
- Money
- Ship
- Trade
- Wheat

B.8 20newsgroups dataset

This well-known repository is composed of 20.000 documents on 20 different topics. This dataset can be downloaded from the UCI Knowledge Discovery in Databases Archive (<http://kdd.ics.uci.edu/>).

Since this repository has the same characteristics as the previous one, it has been adapted to make it suitable for the experiments using the method described for the Reuters dataset.

- alt.atheism
- comp.graphics
- comp.os.ms-windows.misc
- comp.sys.ibm.pc.hardware
- comp.sys.mac.hardware

Argentine grain producers adjusted their yield estimates for the 1986/87 coarse grain crop downward in the week to yesterday after the heavy rains at the end of March and beginning of April, trade sources said. They said sunflower, maize and sorghum production estimates had been reduced despite some later warm, dry weather, which has allowed a return to harvesting in some areas. However, as showers fell intermittently after last weekend, producers feared another spell of prolonged and intense rain could cause more damage to crops already badly hit this season. Rains in the middle of last week reached an average of 27 millimetres in parts of Buenos Aires province, 83 mm in Cordoba, 41 in Santa Fe, 50 in Entre Rios and Misiones, 95 in Corrientes, eight in Chaco and 35 in Formosa. There was no rainfall in the same period in La Pampa. Producers feared continued damp conditions could produce rotting and lead to still lower yield estimates for all the crops, including soybean.

Figure B.7: Reuters. Extract from a document on wheat.

- comp.windows.x
- misc.forsale
- rec.autos
- rec.motorcycles
- rec.sport.baseball
- rec.sport.hockey
- sci.crypt
- sci.electronics
- sci.med
- sci.space
- soc.religion.christian
- talk.politics.guns
- talk.politics.mideast
- talk.politics.misc
- talk.religion.misc

VIKING 1 was launched from Cape Canaveral, Florida on August 20, 1975 on a TITAN 3E-CENTAUR D1 rocket. The probe went into Martian orbit on June 19, 1976, and the lander set down on the western slopes of Chryse Planitia on July 20, 1976. It soon began its programmed search for Martian micro-organisms (there is still debate as to whether the probes found life there or not), and sent back incredible color panoramas of its surroundings. One thing scientists learned was that Mars' sky was pinkish in color, not dark blue as they originally thought (the sky is pink due to sunlight reflecting off the reddish dust particles in the thin atmosphere). The lander set down among a field of red sand and boulders stretching out as far as its cameras could image.

Figure B.8: 20newsgroups. Extract from a document on *sci.space*.

Appendix C

Queries

This Appendix describes the queries used in the experiments carried out in Chapter 7. Three datasets are used in the experiments developed in that chapter of the thesis. They are the UCM dataset, the *20newsgroups* dataset, and the Reuters dataset. Although all of them are described in depth in Appendix B, this section summarizes their main characteristics.

- **UCM dataset:**
 - Number of documents: 104.
 - Number of topics: 11.
 - Kind of documents: Papers.
- ***20newsgroups* dataset:**
 - Number of documents: 20000.
 - Number of topics: 20.
 - Kind of documents: One file per topic (adaptation described in Section B.7).
- **Reuters dataset:**
 - Number of documents: 200.
 - Number of topics: 10.
 - Kind of documents: One file per topic (adaptation described in Section B.7).

Similarly, the main characteristics of the queries used in the experiments carried out in Chapter 7 are as follows:

- **UCM dataset:**
 - Kind of queries: Abstracts.
 - Number of queries: 4.
 - Sizes of queries:
 - * 2 x 1KB.
 - * 2 x 2KB.
- ***20newsgroups* dataset:**
 - Kind of queries: Messages.
 - Number of queries: 10.
 - Sizes of queries:
 - * 7 x 2KB.
 - * 3 x 3KB.
- **Reuters dataset:**
 - Kind of queries: News.
 - Number of queries: 10.
 - Sizes of queries:
 - * 10 x 2KB.

Agent-based modelling facilitates the implementation of tools for the analysis of social patterns. This comes from the fact that agent related concepts allow the representation of organizational and behavioural aspects of individuals in a society and their interactions. An agent can characterize an individual with capabilities to perceive and react to events in the environment, taking into account its mental state (beliefs, goals), and to interact with other agents in its social environment. There are already tools to perform agent-based social simulation but these are usually hard to use by social scientists, as they require a good expertise in computer programming. In order to cope with such difficulty, we propose the use of agent-based graphical modelling languages, which can help to specify social systems as multi-agent systems in a more convenient way. This is complemented with transformation tools to be able to analyse and derive emergent social behavioural patterns by using the capabilities of existing simulation platforms. In this way, this framework can facilitate the specification and analysis of complex behavioural patterns that may emerge in social systems.

Figure C.1: UCM. Example of query.

As the subjects says, Windows 3.1 keeps crashing (giving me GPF) on me of late. It was never a very stable package, but now it seems to crash every day. The worst part about it is that it does not crash consistently: ie I. There is a way in SYS.INI to turn off RAM parity checking (unfortunately, my good Windows references are at home, but any standard Win reference will tell you how to do it. If not, email back to me). That weird memory may be producing phony parity errors. Danger is, if you turn checking off, you run the slight risk of data corruption due to a missed real error. I had this very same problem, and did 'work around' by turning parity checking off, but that only worked while I was in windows, and the parity error would occur immediately after exiting windows, however, the problem turned out to be 3 chip simms vs 9 chip simms. I can't use 3 chip simms in my computer, and when I replaced them, the problem vanished, forever.

Figure C.2: 20newsgroups. Example of query.

FAO SEES LOWER GLOBAL WHEAT, GRAIN OUTPUT IN 1987.
The U.N Food and Agriculture Organisation (FAO) said global wheat and coarse grain output was likely to fall in 1987 but supplies would remain adequate to meet demand. FAO said in its monthly food outlook bulletin total world grain output was expected to fall 38 mln tonnes to 1,353 mln in 1987, due mainly to unusually high winter losses in the Soviet Union, drought in China and reduced plantings in North America. World cereal stocks at the end of 1986/87 were forecast to rise 47 mln tonnes to a record 452 mln tonnes, softening the impact of reduced production. But stocks are unevenly distributed, with about 50 pct held by the U.S. "Thus the food security prospects in 1987/88 for many developing countries, particularly in Africa, depend crucially on the outcome of this year's harvests", FAO said. FAO said world cereal supplies in 1986/87 were estimated at a record 2,113 mln tonnes, about five pct higher than last season and due mainly to large stocks and a record 1986 harvest, estimated at 1,865 mln tonnes. FAO's forecast of 1986/87 world cereals trade was revised upwards by eight mln tonnes to 179 mln due to the likelihood of substantial buying by China and the Soviet Union.

Figure C.3: Reuters. Example of query.

Appendix D

Detailed Experimental Results

D.1 Preliminary study on text distortion

This section shows all the results obtained in the work developed in Chapter 5. Three clustering error figures are shown for each *dataset-compression algorithm* pair. Each of them corresponds to a selection method:

- *MFW selection method*
- *RW selection method*
- *LFW selection method*

In all the clustering error figures, the value on the horizontal axis corresponds to the cumulative sum of the BNC-based frequencies of the words substituted from the documents, whereas the value on the vertical axis corresponds to the clustering error. Furthermore, the curves with asterisk markers correspond to the *asterisk substitution method*, while the curves with square markers correspond to the *random character substitution method*.

Analyzing all the figures one can observe that the *asterisk substitution method* is always better than the *random character substitution method*. This is to be expected because substituting a word with random characters adds noise to the documents, and therefore most likely increases the Kolmogorov complexity of the documents and makes the clustering worse. In addition, it can be seen that the best clustering results correspond to the *MFW selection method*, the worst results correspond to the *LFW selection method*, and the results corresponding to the *RW selection method* are situated in between them.

The dendrogram obtained with no distortion, and the best dendrogram obtained are shown for each *dataset-compression algorithm* pair as well. In the cases in which the dendrogram obtained with no distortion is the best one obtained, only one dendrogram is shown.




Figure D.1: Books. PPMZ compressor. MFW selection method.




Figure D.2: Books. PPMZ compressor. RW selection method.




Figure D.3: Books. PPMZ compressor. LFW selection method.

The combination of the *MFW selection method* and the *asterisk substitution method* improves the clustering results so much that a clustering error of 0 is obtained when the texts are distorted using the set of words that accumulate a BNC-based frequency of 0.9.

- Non-distorted clustering error: 5
- Best clustering error: 0

The improvement can be observed by comparing Figs D.4 and D.5. Whereas the books by Edgar Allan Poe and Alexander Pope are not correctly clustered in Fig D.4, all the books are correctly clustered in Fig D.5. That is the reason why the clustering error that corresponds to the best dendrogram obtained is 0.

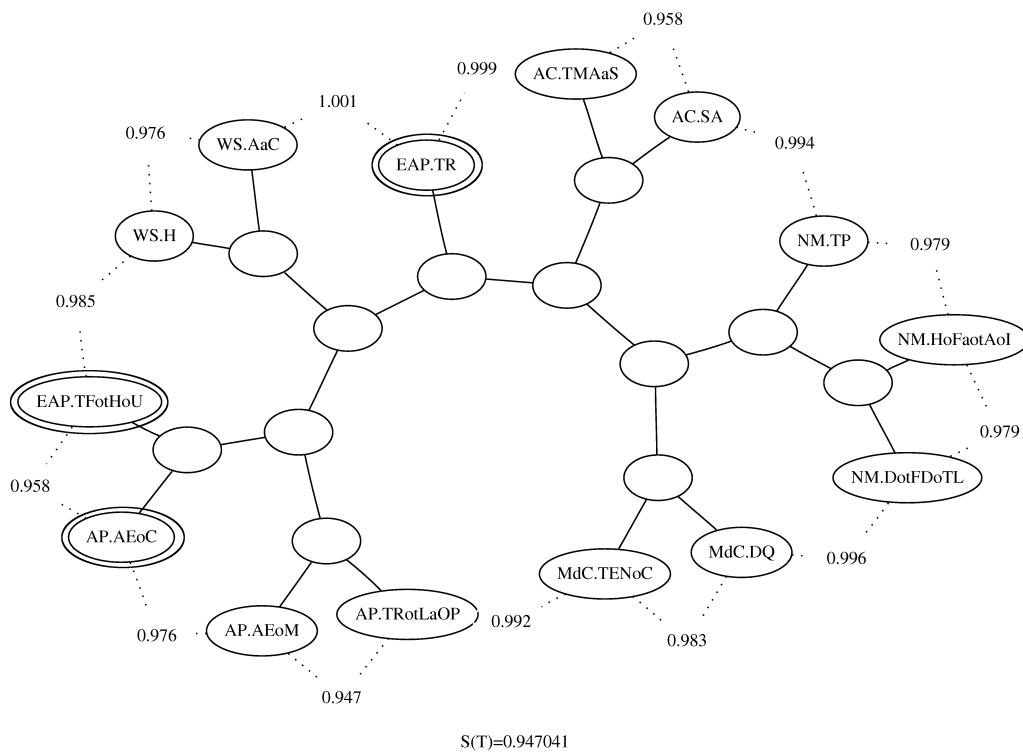


Figure D.4: Books. PPMZ compressor. Dendrogram obtained with no distortion.




Figure D.6: Books. LZMA compressor. MFW selection method.




Figure D.7: Books. LZMA compressor. RW selection method.




Figure D.8: Books. LZMA compressor. LFW selection method.

The combination of the *MFW selection method* and the *asterisk substitution method* improves the clustering results when the texts are distorted using the sets of words that accumulate a BNC-based frequency of 0.8 and 0.9.

- Non-distorted clustering error: 4
- Best clustering error: 2

The improvement can be observed by comparing Figs D.9 and D.10. The difference between both figures is that the book “The Prince” by Niccolò Machiavelli is closer to the rest of Niccolò Machiavelli’s books in Fig D.10.

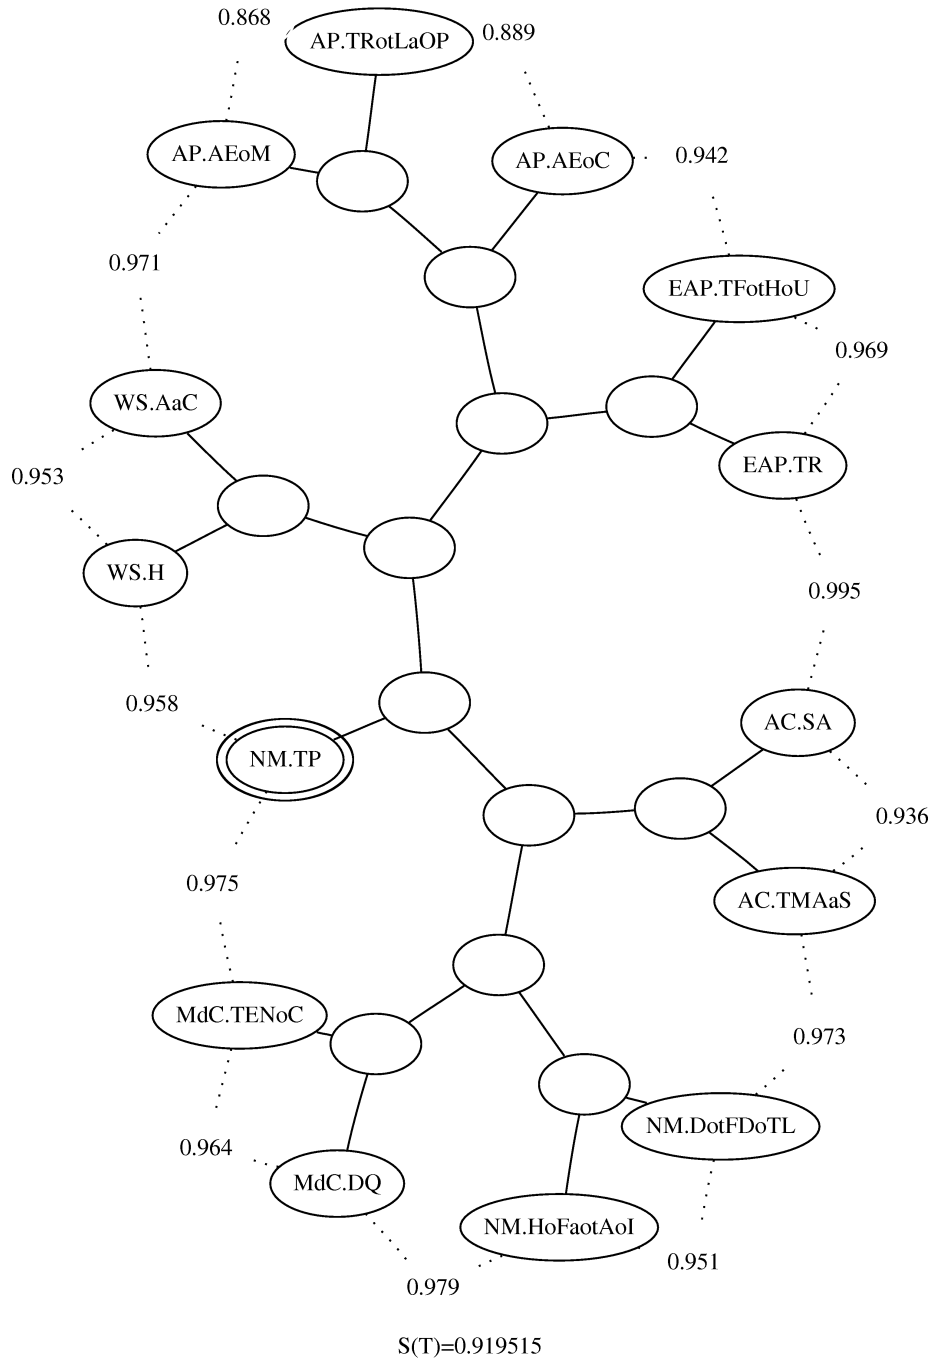


Figure D.9: Books. LZMA compressor. Dendrogram obtained with no distortion.

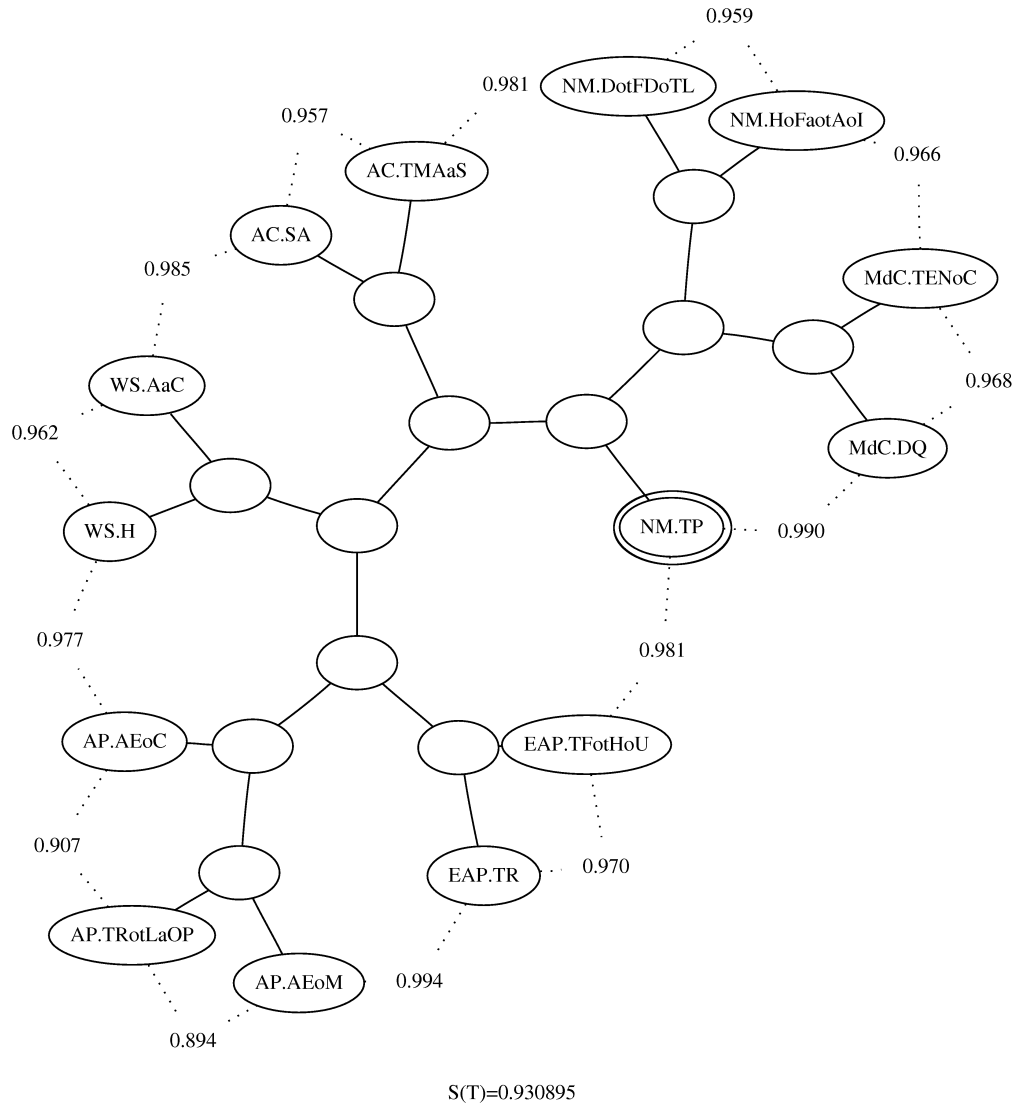


Figure D.10: Books. LZMA compressor. Best dendrogram obtained.




Figure D.11: Books. BZIP2 compressor. MFW selection method.




Figure D.12: Books. BZIP2 compressor. RW selection method.




Figure D.13: Books. BZIP2 compressor. LFW selection method.

Again, the combination of the *MFW selection method* and the *asterisk substitution method* improves the clustering results when the texts are distorted using the sets of words that accumulate a BNC-based frequency of 0.7 and 0.8. However, in this case the non-distorted clustering error is improved using the rest of the selection methods, as can be observed looking at Figs D.12 and D.13.

The problematic books in this case are the books by Miguel de Cervantes and the books by Niccolò Machiavelli. The difference between the dendrogram obtained with no distortion, and the best dendrogram obtained is that the book “The Prince” by Niccolò Machiavelli is closer to the rest of Niccolò Machiavelli’s books in Fig D.15.

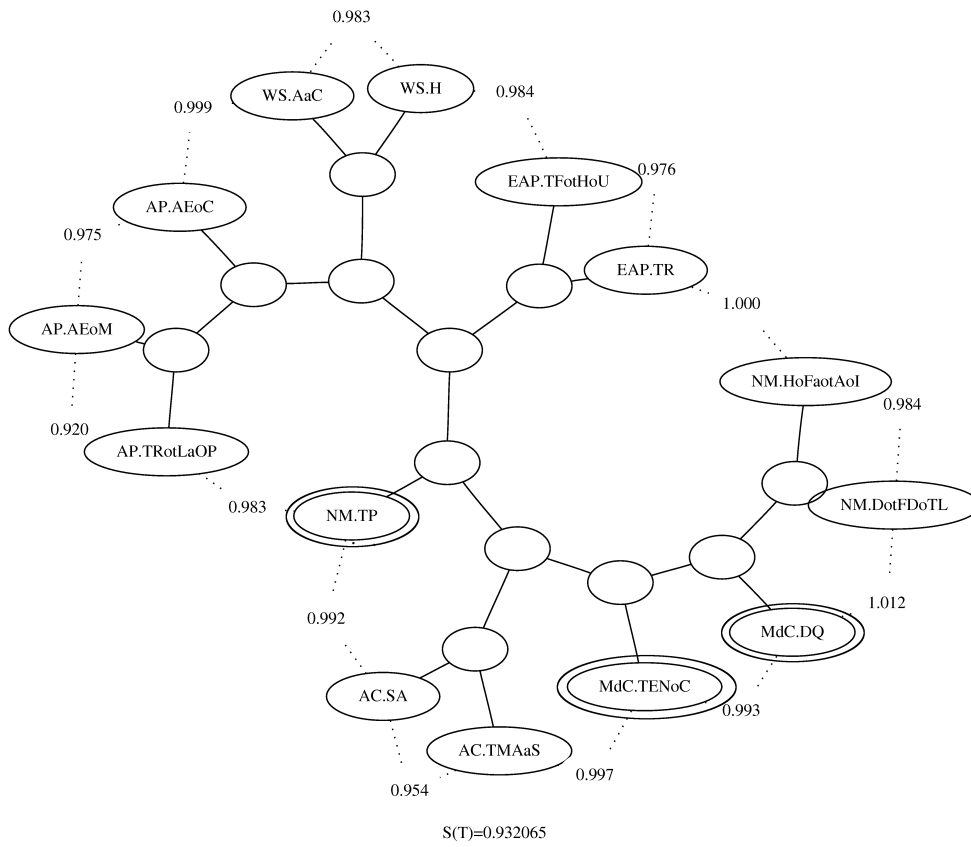


Figure D.14: Books. BZIP2 compressor. Dendrogram obtained with no distortion.

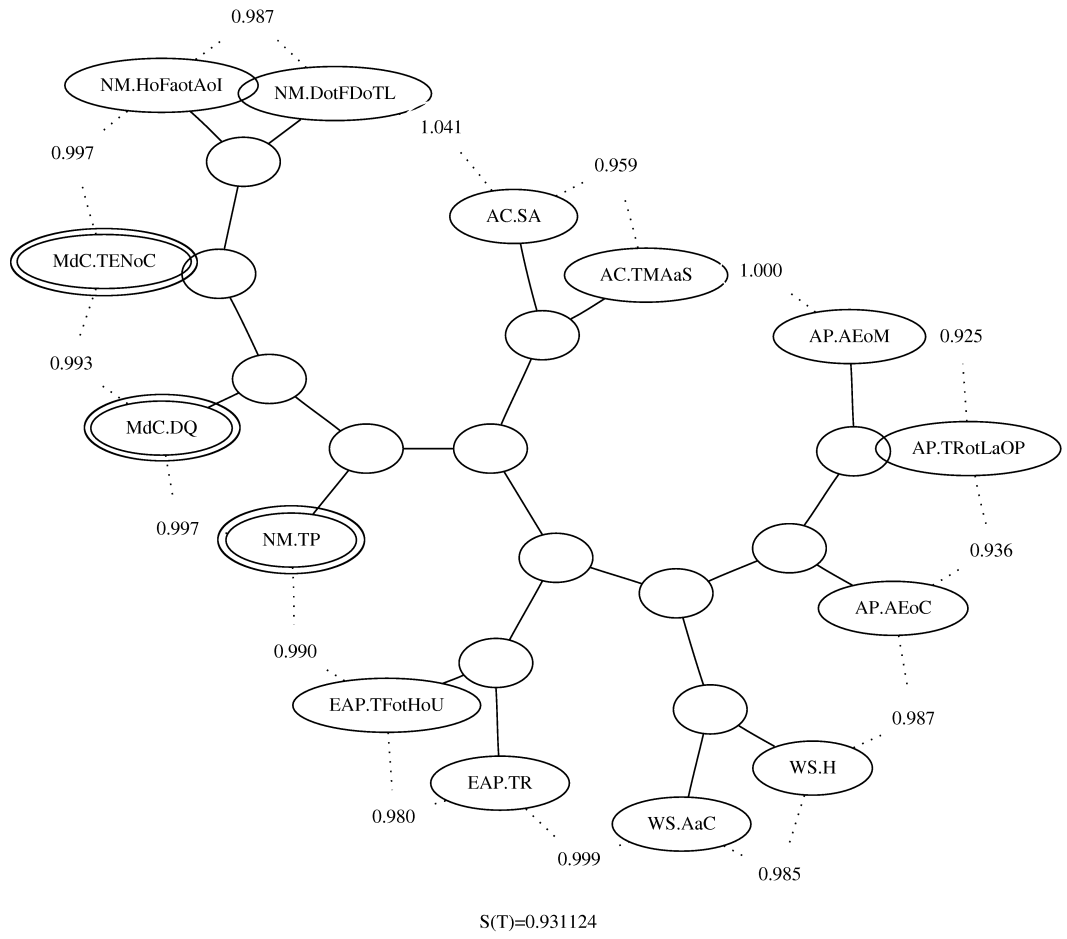


Figure D.15: Books. BZIP2 compressor. Best dendrogram obtained.




Figure D.16: UCI-KDD. PPMZ compressor. MFW selection method.




Figure D.17: UCI-KDD. PPMZ compressor. RW selection method.




Figure D.18: UCI-KDD. PPMZ compressor. LFW selection method.

The non-distorted clustering error in this case is 0. Therefore, it is impossible to improve the clustering error for this *dataset-compression algorithm* pair.

However, analyzing Fig D.16 one can observe that the clustering error remains constant from the point that corresponds to a BNC-based frequency of 0 to the one corresponding to a BNC-based frequency of 0.8. This means that the relevant information contained in the documents is maintained despite the word removal.

Again, the results show that the combination of the *MFW selection method* and the *asterisk substitution method* is the key factor, because whereas a clustering error of 0 is obtained from most of the points of the curve with asterisk markers in Fig D.16, the clustering error in the other cases gets worse, as one can see in Figs D.17, and D.18.

Since the dendrogram obtained with no distortion clusters all the texts perfectly, only one dendrogram is shown for this *dataset-compression algorithm* pair.

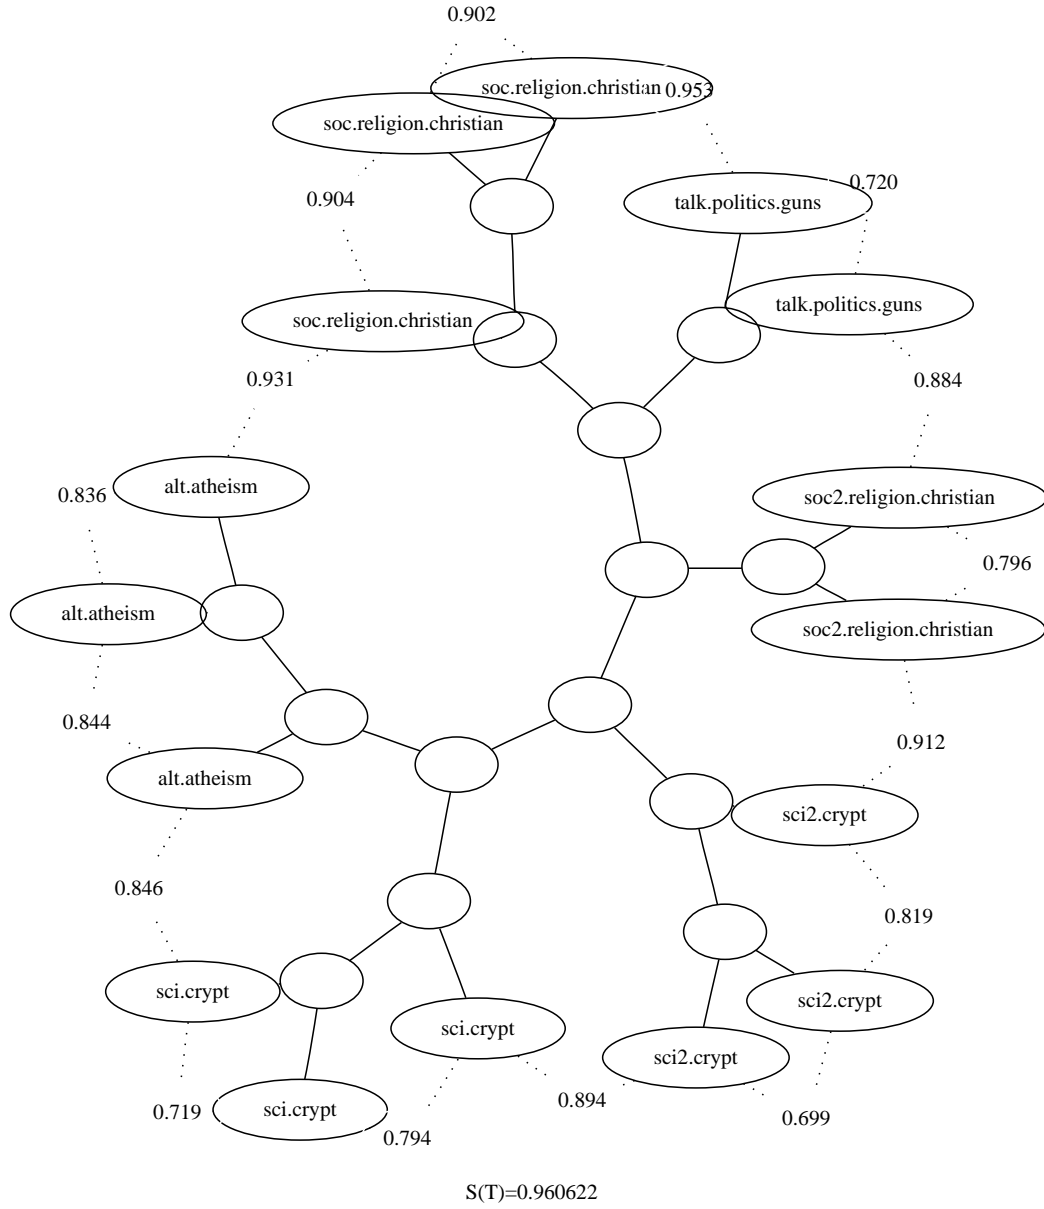


Figure D.19: UCI-KDD. PPMZ compressor. Best dendrogram obtained.




Figure D.20: UCI-KDD. LZMA compressor. MFW selection method.




Figure D.21: UCI-KDD. LZMA compressor. RW selection method.




Figure D.22: UCI-KDD. LZMA compressor. LFW selection method.

Similarly to the results shown before, the non-distorted clustering error in this case is 0. Therefore, it is impossible to improve the clustering error for this *dataset-compression algorithm* pair.

However, analyzing Fig D.20 it can be observed that the clustering error remains constant when the *MFV selection method* and the *asterisk substitution method* are combined. This behavior is observed for the points from 0.0 to 0.9 of the curve.

In this case, the results that correspond to the rest of selection methods are almost the same as the ones obtained using the *MFV selection method*.

Since the dendrogram obtained with no distortion clusters all the texts perfectly, only one dendrogram is shown for this *dataset-compression algorithm* pair.




Figure D.24: UCI-KDD. BZIP2 compressor. MFW selection method.




Figure D.25: UCI-KDD. BZIP2 compressor. RW selection method.




Figure D.26: UCI-KDD. BZIP2 compressor. LFW selection method.

Again, the non-distorted clustering error in this case is 0. Therefore, it is impossible to improve the clustering error for this *dataset-compression algorithm* pair.

However, analyzing Fig D.24 it can be observed that the clustering error remains constant from a distortion of 0 to a distortion of 0.6.

In this case, the results show that the combination of the *MFW selection method* and the *asterisk substitution method* is the key factor, because whereas a clustering error of 0 is obtained from most of the points of the curve with asterisk markers in Fig D.24, the clustering error in the other cases gets worse, as one can see in Figs D.25, and D.26.

Since the dendrogram obtained with no distortion clusters all the texts perfectly, only one dendrogram is shown for this *dataset-compression algorithm* pair.




Figure D.27: UCI-KDD. BZIP2 compressor. Best dendrogram obtained.




Figure D.28: MedlinePlus. PPMZ compressor. MFW selection method.




Figure D.29: MedlinePlus. PPMZ compressor. RW selection method.




Figure D.30: MedlinePlus. PPMZ compressor. LFW selection method.

The combination of the *MFV selection method* and the *asterisk substitution method* improves the clustering results as follows:

- Non-distorted clustering error: 14
- Best clustering error: 4

This improvement can be observed by comparing Figs D.31 and D.32. Whereas three documents are not correctly clustered in Fig D.31, only two documents are not correctly clustered in Fig D.32. Furthermore, in the best dendrogram obtained, the documents that are not correctly clustered are adjacent to the ones related to them.

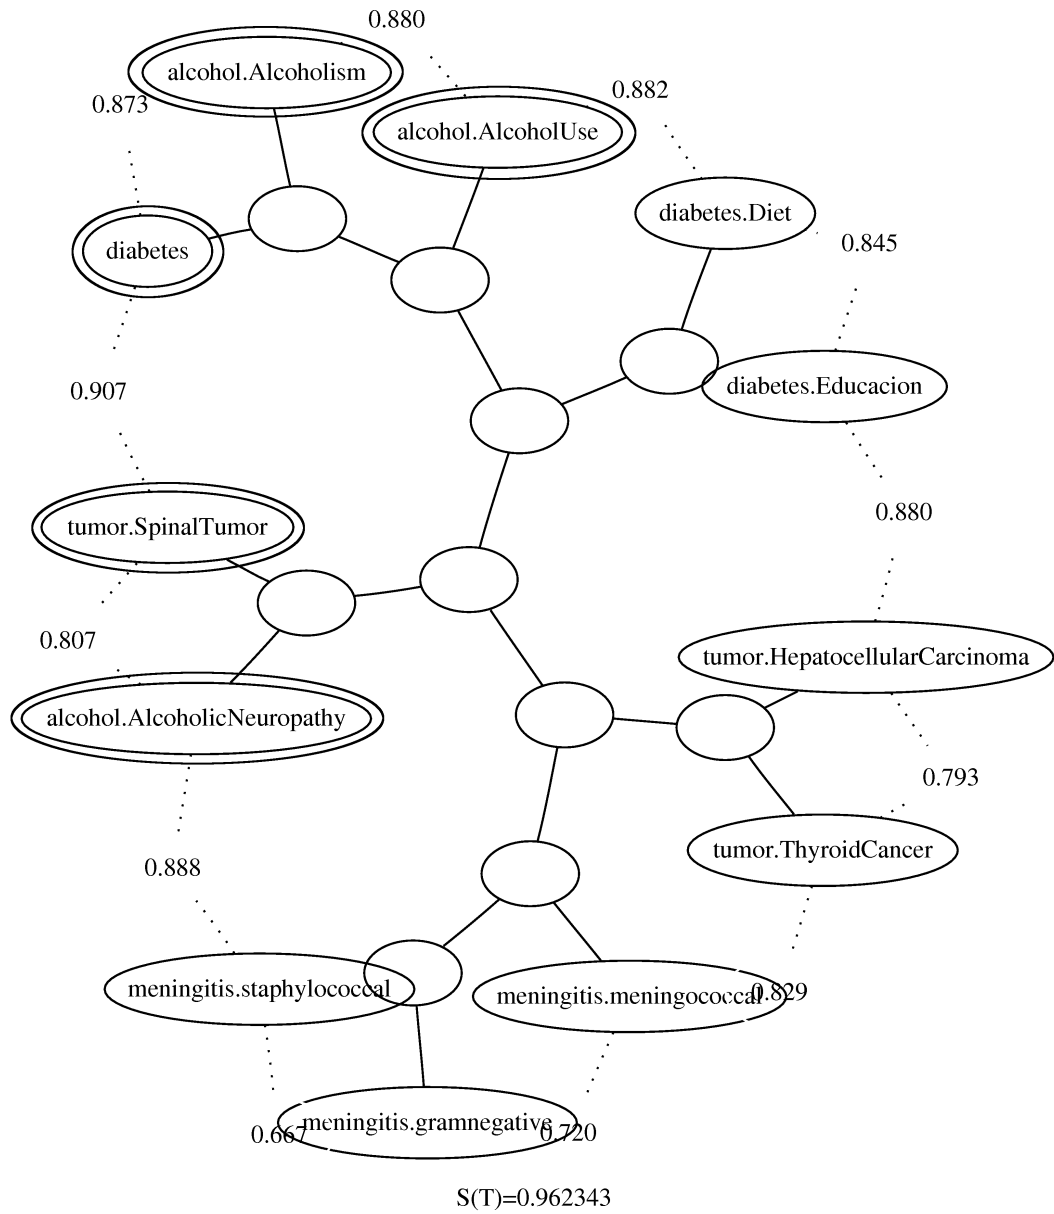


Figure D.31: MedlinePlus. PPMZ compressor. Dendrogram obtained with no distortion.

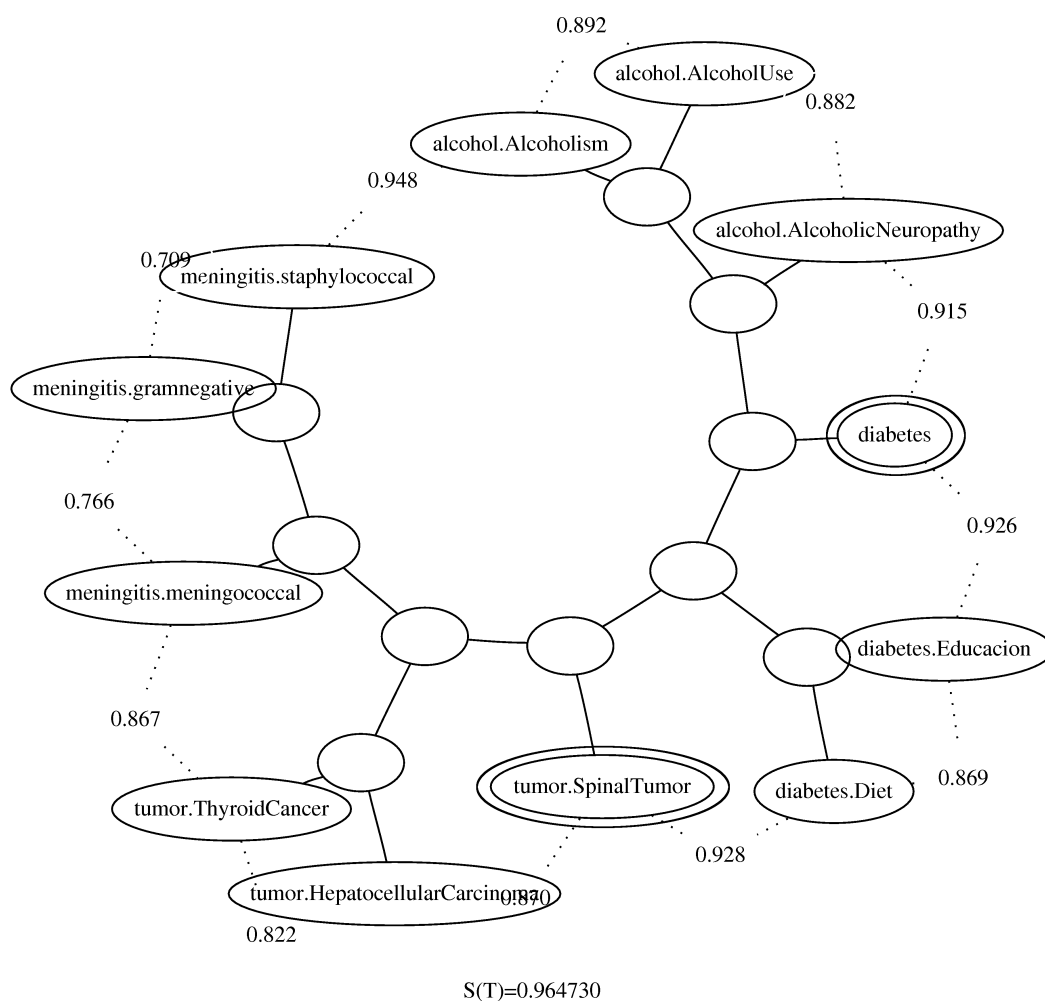


Figure D.32: MedlinePlus. PPMZ compressor. Best dendrogram obtained.




Figure D.33: MedlinePlus. LZMA compressor. MFW selection method.




Figure D.34: MedlinePlus. LZMA compressor. RW selection method.




Figure D.35: MedlinePlus. LZMA compressor. LFW selection method.

The combination of the *MFV selection method* and the *asterisk substitution method* improves the clustering results when the texts are distorted using the set of words that accumulate a BNC-based frequency of 0.7 and 0.8.

- Non-distorted clustering error: 14
- Best clustering error: 10

As usual, the improvement can be observed by comparing Figs D.36 and D.37. Three texts are problematic in this case. They are about diabetes, tumor and alcohol. The difference between the dendrogram obtained with no distortion and the best dendrogram obtained is that, in the latter, the texts are closer to the texts that are related to them.

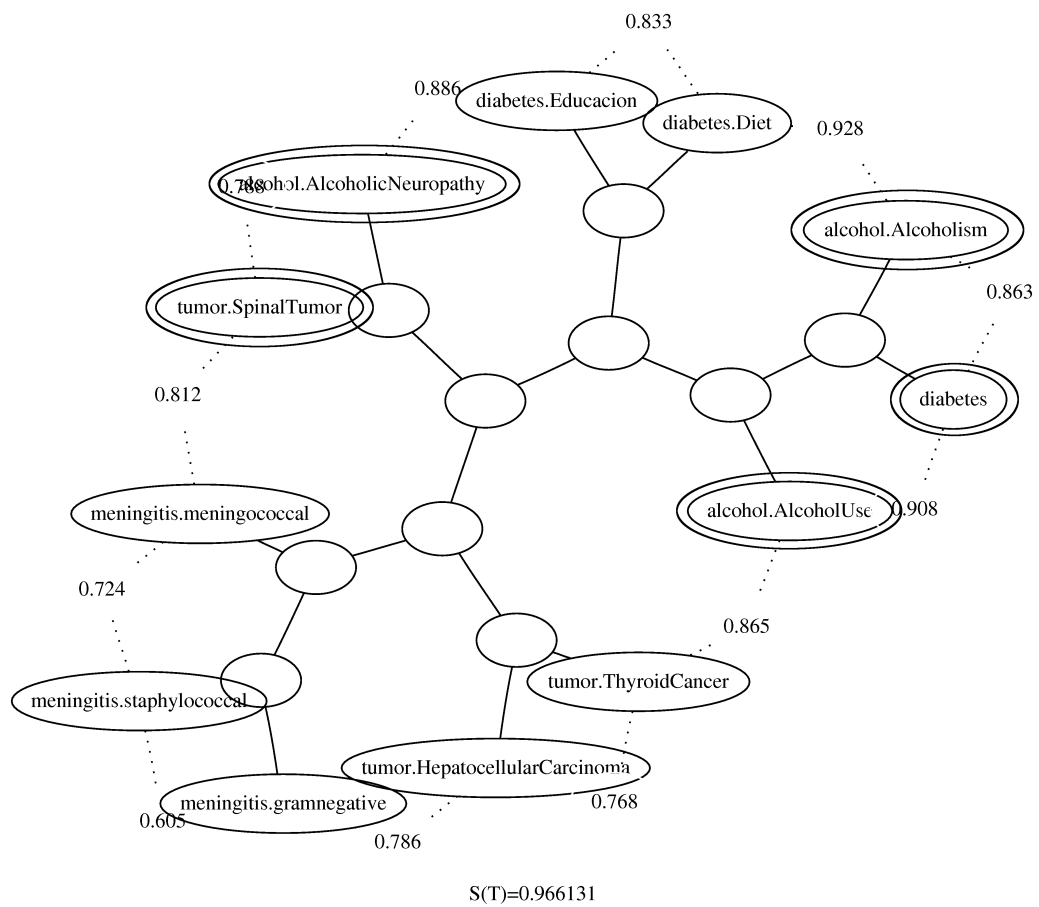


Figure D.36: MedlinePlus. LZMA compressor. Dendrogram obtained with no distortion.

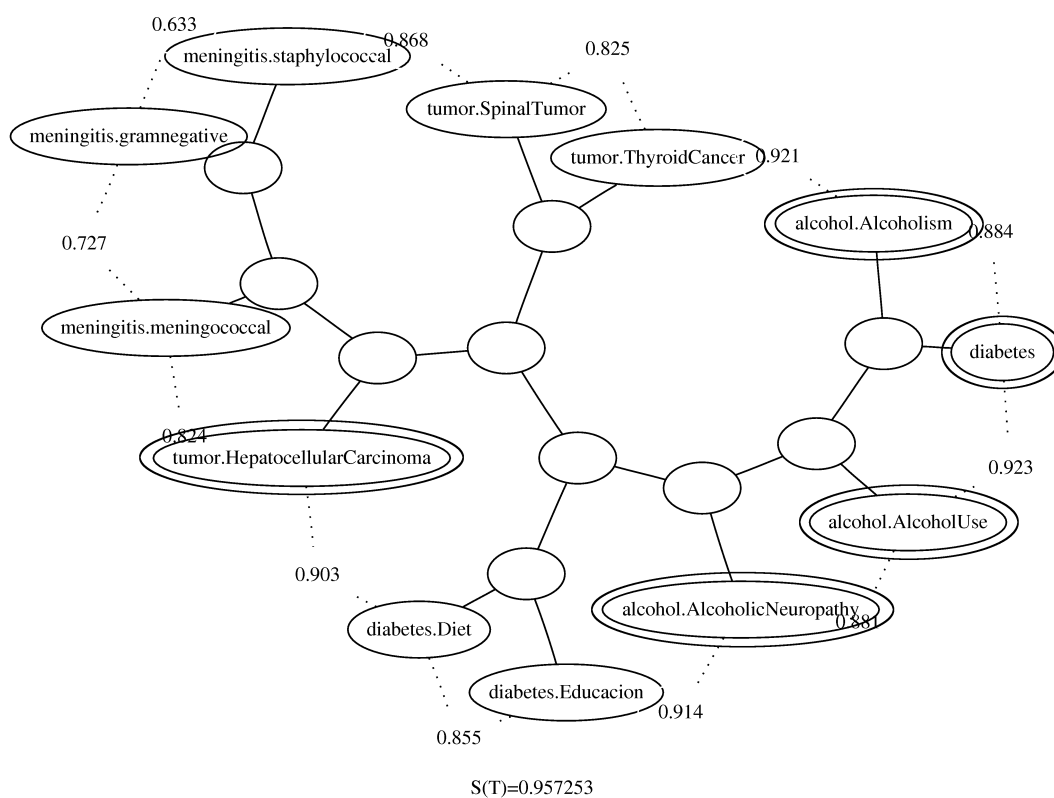


Figure D.37: MedlinePlus. LZMA compressor. Best dendrogram obtained.




Figure D.38: MedlinePlus. BZIP2 compressor. MFW selection method.




Figure D.39: MedlinePlus. BZIP2 compressor. RW selection method.




Figure D.40: MedlinePlus. BZIP2 compressor. LFW selection method.

Analyzing Figs D.38, D.39, and D.40 one can observe that the best clustering results correspond to the *MFW selection method*, the worst results correspond to the *LFW selection method*, and the results corresponding to the *RW selection method* are situated in between them.

As usual, the combination of the *MFW selection method* and the *asterisk substitution method* improves the clustering results. In this case, this improvement is obtained when the texts are distorted using the set of words that accumulate a BNC-based frequency of 0.5, 0.7 and 0.8.

- Non-distorted clustering error: 14
- Best clustering error: 10

Again, comparing Figs D.36 and D.37 one can see the clustering behavior improvement. For this *dataset-compression algorithm* pair three texts are problematic. They are about diabetes, tumor and alcohol. The difference between the dendrogram obtained with no distortion and the best dendrogram obtained is that in the latter the texts are closer to the texts to which they are related.

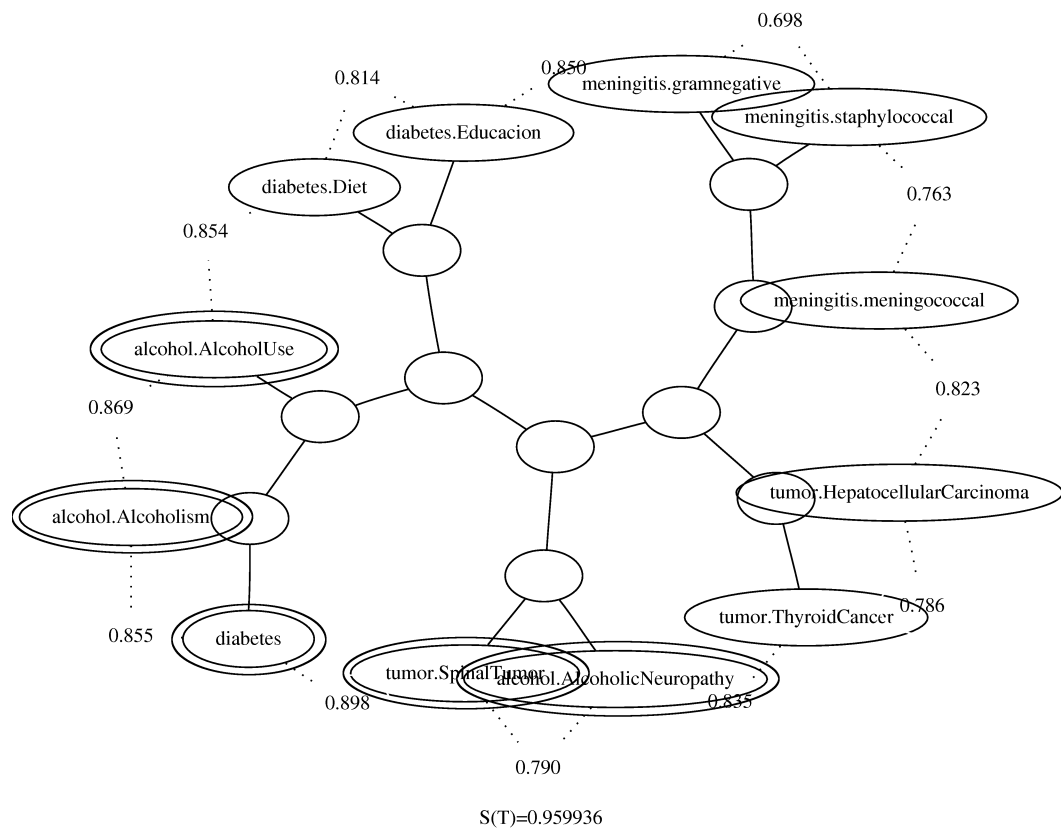


Figure D.41: MedlinePlus. BZIP2 compressor. Dendrogram obtained with no distortion.

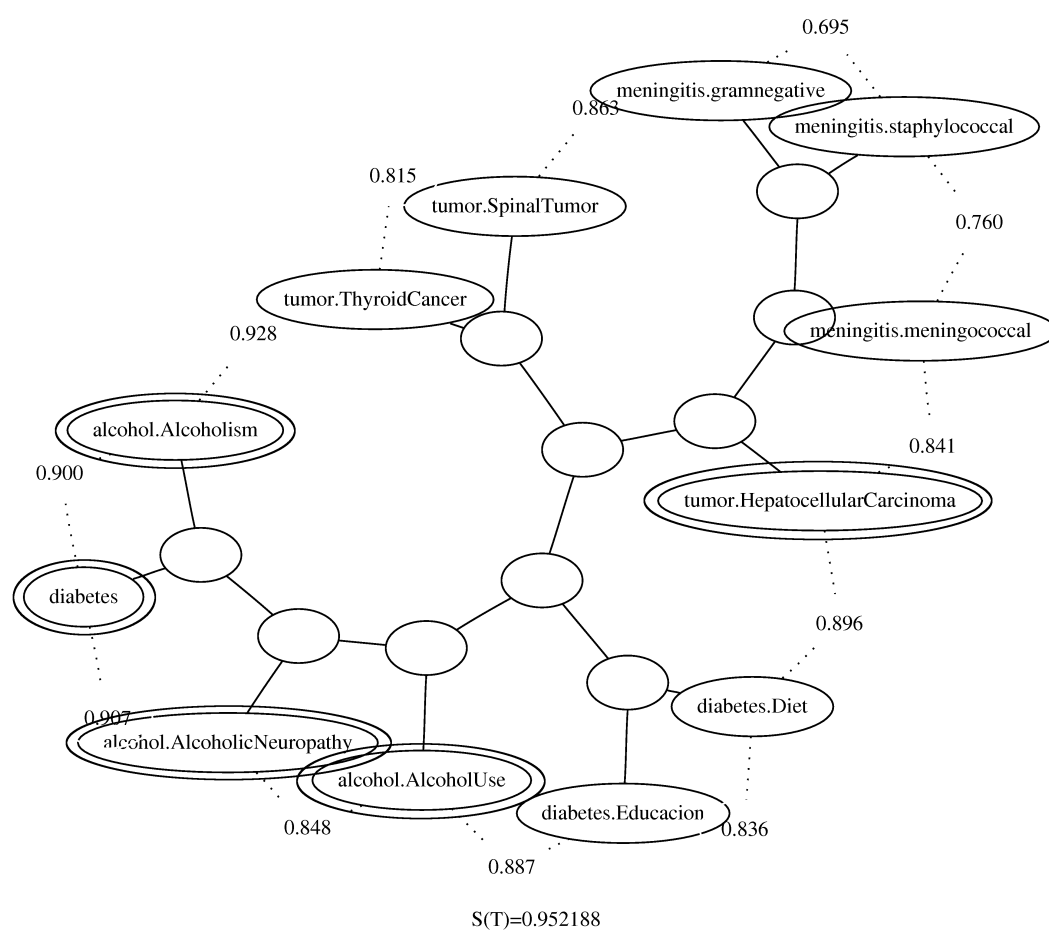


Figure D.42: MedlinePlus. BZIP2 compressor. Best dendrogram obtained.




Figure D.43: IMDB. PPMZ compressor. MFW selection method.




Figure D.44: IMDB. PPMZ compressor. RW selection method.




Figure D.45: IMDB. PPMZ compressor. LFW selection method.

Again, analyzing Figs D.43, D.44, and D.45 one can observe that the best clustering results correspond to the *MFW selection method*, the worst results correspond to the *LFW selection method*, and the results corresponding to the *RW selection method* are situated in between them.

Given that the dendrogram obtained with no distortion clusters all the texts perfectly, only one dendrogram is shown in this case.

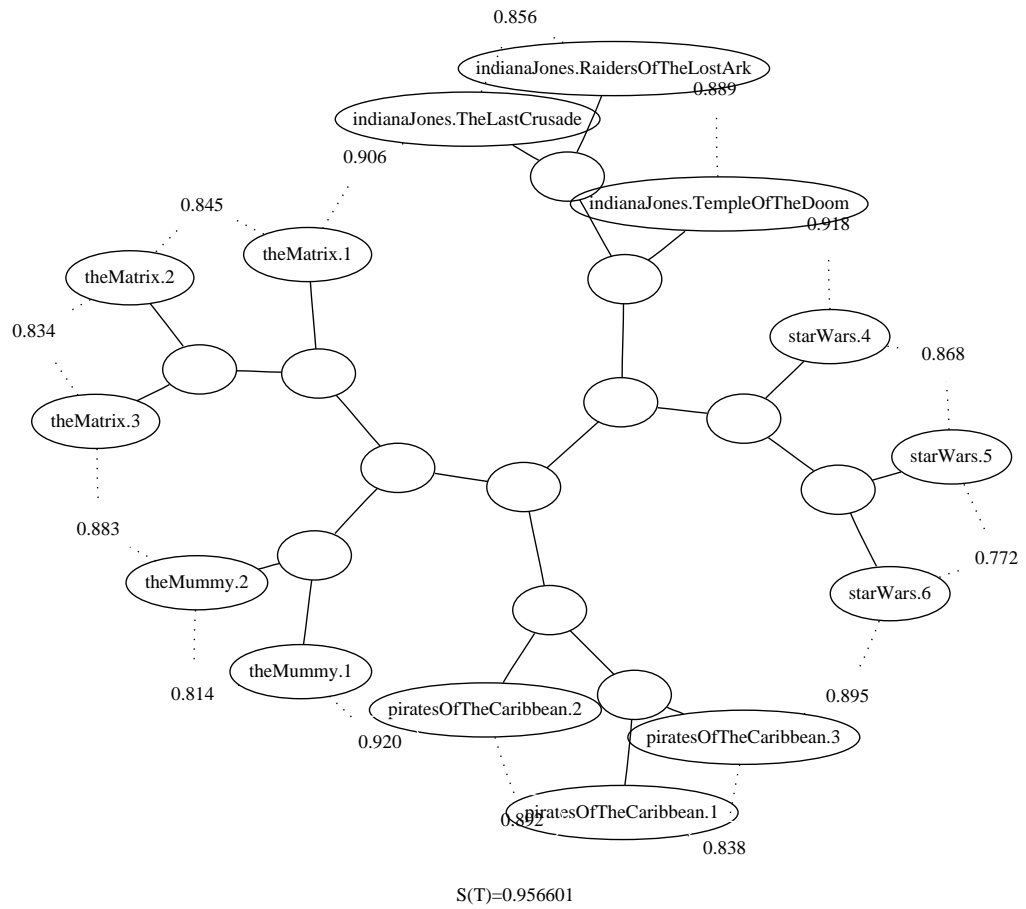


Figure D.46: IMDB. PPMZ compressor. Best dendrogram obtained.




Figure D.47: IMDB. LZMA compressor. MFW selection method.




Figure D.48: IMDB. LZMA compressor. RW selection method.




Figure D.49: IMDB. LZMA compressor. LFW selection method.

The combination of the *MFW selection method* and the *asterisk substitution method* improves the clustering results when the texts are distorted using the sets of words that accumulate a BNC-based frequency of 0.6, 0.7, 0.8 and 0.9.

- Non-distorted clustering error: 18
- Best clustering error: 4

The improvement can be observed by comparing Figs D.50 and D.51. The difference between both figures is that the movies “Pirates of the Caribbean 2”, “Star Wars 4” and “Indiana Jones and the Temple of the Doom” are incorrectly clustered in the dendrogram depicted in Fig D.50, whereas only the movie “Pirates of the Caribbean 2” is incorrectly clustered in Fig D.51.

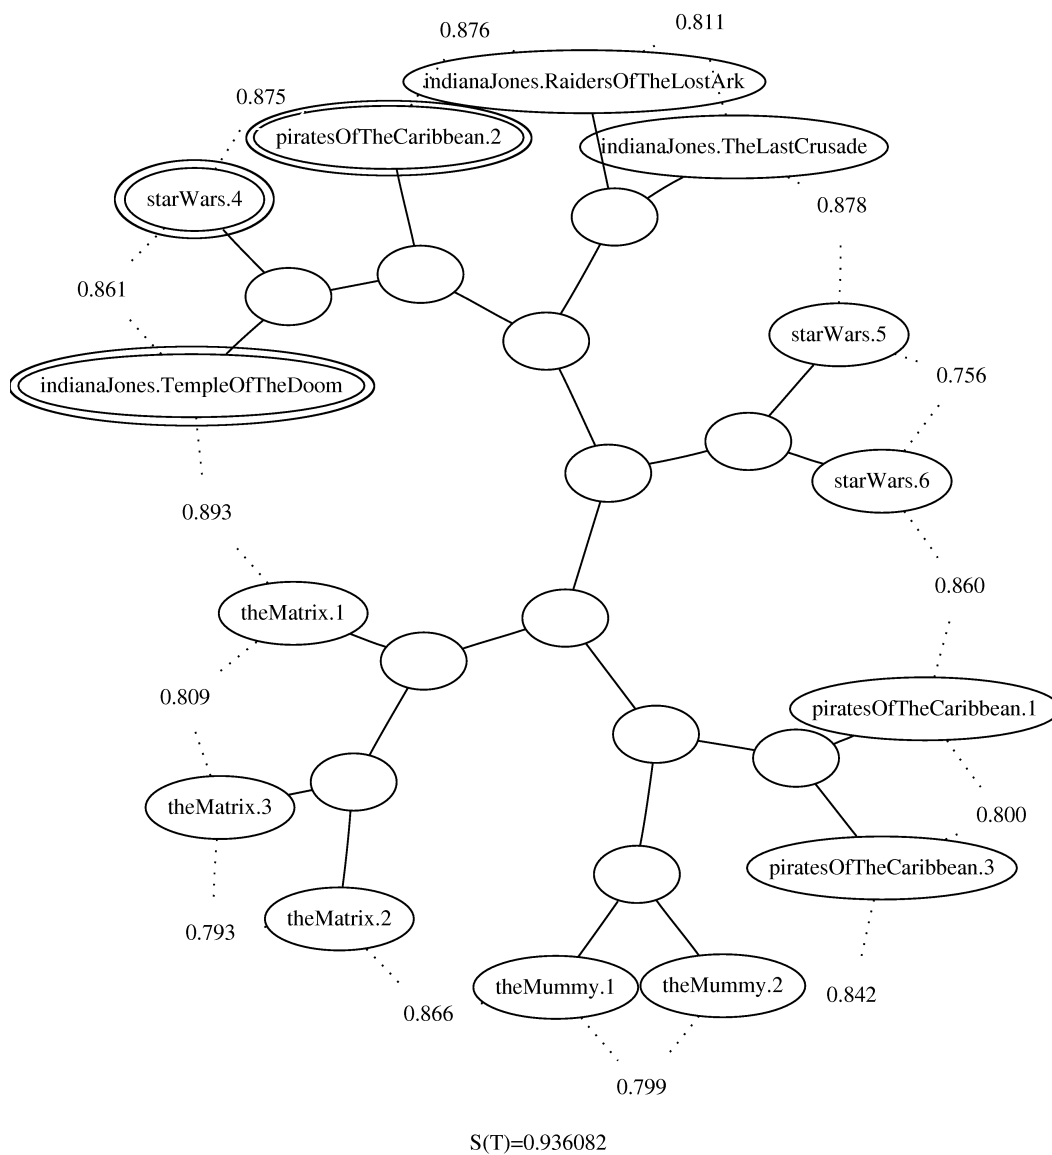


Figure D.50: IMDB. LZMA compressor. Dendrogram obtained with no distortion.

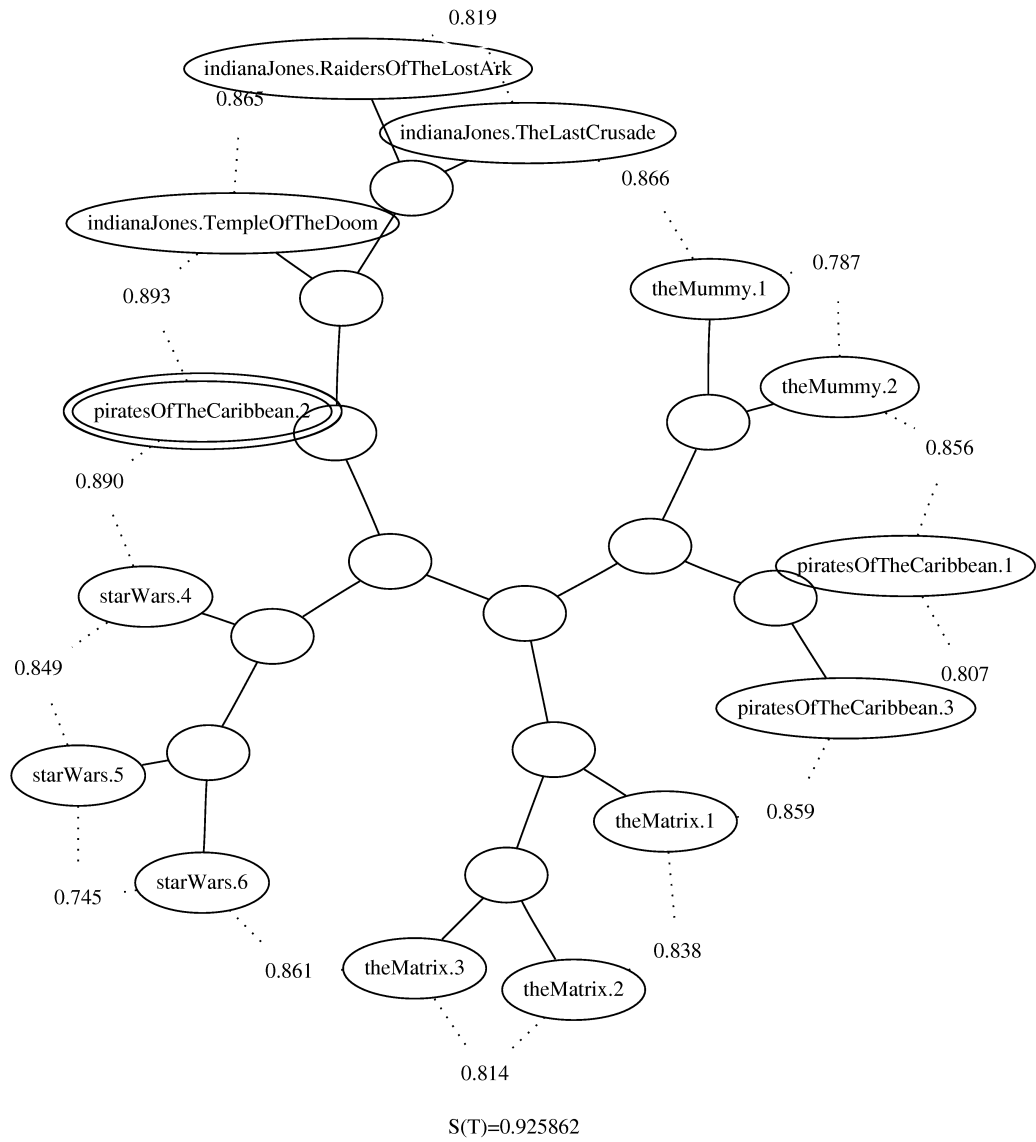


Figure D.51: IMDB. LZMA compressor. Best dendrogram obtained.




Figure D.52: IMDB. BZIP2 compressor. MFW selection method.




Figure D.53: IMDB. BZIP2 compressor. RW selection method.




Figure D.54: IMDB. BZIP2 compressor. LFW selection method.

The clustering error obtained without distortion is 0, therefore, only one dendrogram is depicted for this *dataset-compression algorithm* pair.

Again, analyzing Figs D.52, D.53, and D.54, one can reach the conclusion that the best clustering results correspond to the *MFW selection method*, the worst results correspond to the *LFW selection method*, and the results corresponding to the *RW selection method* are situated in between them.

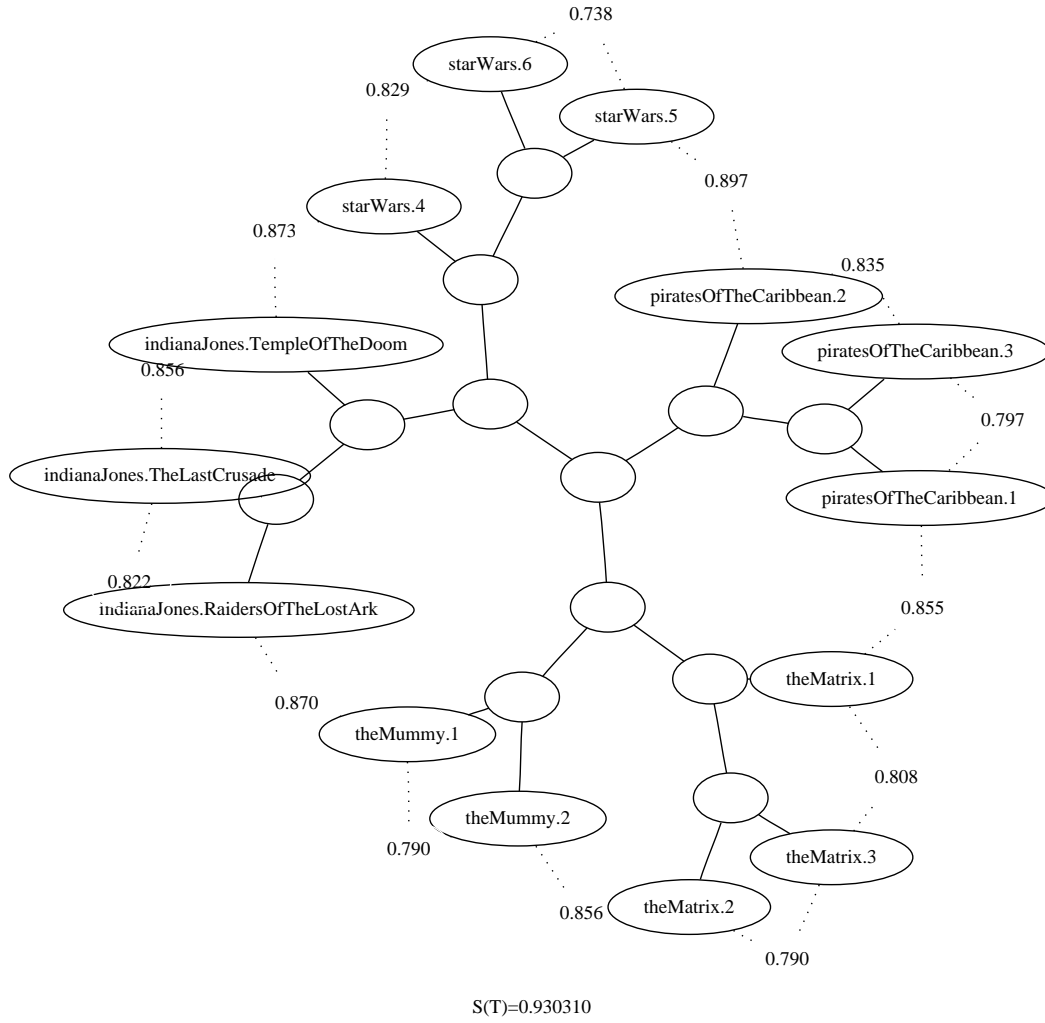


Figure D.55: IMDB. BZIP2 compressor. Best dendrogram obtained.

Bibliography

- [1] O. Acosta, J. Fripp, A. Rueda, D. Xiao, E. Bonner, P. Bourgeat, and O. Salvado. 3D Shape Context Surface Registration for Cortical Mapping. In *Proceedings of the IEEE International Symposium on Biomedical Imaging*, ISBI '10, pages 1021–1024, 2010.
- [2] G. Adomavicius, R. Sankaranarayanan, S. Sen, and A. Tuzhilin. Incorporating Contextual Information in Recommender Systems using a Multidimensional Approach. *ACM Transactions on Information Systems*, 23(1):103–145, 2005.
- [3] E. B. Allen, T. M. Khoshgoftaar, and Y. Chen. Measuring Coupling and Cohesion of Software Modules: An Information-Theory Approach. In *Proceedings of the 7th International Symposium on Software Metrics*, METRICS '01, pages 124–, 2001.
- [4] M. Apel, C. Bockermann, and M. Meier. Measuring Similarity of Malware Behavior. In *IEEE 34th Conference on Local Computer Networks*, pages 891–898, 2009.
- [5] T. Arbuckle, A. Balaban, D. K. Peters, and M. Lawford. Software Documents: Comparison and Measurement. In *Proceedings of the 18th Int. Conf. on Software Engineering and Knowledge Engineering*, SEKE '07, pages 740–745, 2007.
- [6] R. Baeza-Yates and B. Ribeiro-Neto. *Modern Information Retrieval*. Addison-Wesley, 1999.
- [7] K. J. Balakrishnan and N. A. Toubia. Relating entropy theory to test data compression. In *Proceedings of the European Test Symposium*, pages 94–99, 2004.
- [8] T. Bell, I. H. Witten, and J. G. Cleary. Modeling for Text Compression. *ACM Computing Surveys*, 21(4):557–591, 1987.

- [9] S. Belongie, J. Malik, and J. Puzicha. Shape Matching and Object Recognition Using Shape Contexts. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 24(4):509–522, 2002.
- [10] D. Benedetto, E. Caglioti, and V. Loreto. Language Trees and Zipping. *Physical Review Letters*, 88(4):48702, 2002.
- [11] J. L. Bentley, D. D. Sleator, R. E. Tarjan, and V. K. Wei. A Locally Adaptive Data Compression Scheme. *Communications of the ACM*, 29(4):320–330, 1986.
- [12] M. Bertacchini and P. I. Fierens. Preliminary Results on Masquerader Detection using Compression Based Similarity Metrics. *Electronic Journal of SADIO*, 7(1):31–42, 2007.
- [13] C. Bloom. PPMZ, [Online] Available: <http://www.cbloom.com>.
- [14] C. Bloom. *Deobfuscating LZMA*. [Online] Available: <http://cbloomrants.blogspot.com/2010/08/08-20-10-deobfuscating-lzma.html>.
- [15] C. R. Bloom. *Solving the Problems of Context Modeling*. [Online] Available from <http://www.cbloom.com/papers/ppmz.zip>.
- [16] C. Bouras and V. Tsogkas. Noun Retrieval Effect on Text Summarization and Delivery of Personalized News Articles to the User’s Desktop. *Data & Knowledge Engineering*, 69(7):664–677, 2010.
- [17] British National Corpus, BNC Consortium, Oxford University Computing Services [Online] Available: <http://www.natcorp.ox.ac.uk/>.
- [18] P. F. Brown, P. V. de Souza, R. L. Mercer, V. J. D. Pietra, and J. C. Lai. Class-based n-gram Models of Natural Language. *Computational Linguistics*, 18(4):467–479, 1992.
- [19] M. Burrows and D. Wheeler. A block-sorting lossless data compression algorithm. Technical report, 1994.
- [20] J. P. Callan. Passage-Level Evidence in Document Retrieval. In *Proceedings of the 17th annual international ACM SIGIR conference on Research and development in information retrieval*, SIGIR ’94, pages 302–310, 1994.

- [21] T. Caus, S. Christmann, and S. Hagenhoff. Development of Context-Aware Mobile Services: an Approach to Simplification. *International Journal of Mobile Communications*, 7(2):133–153, 2009.
- [22] M. Cebrián, M. Alfonseca, and A. Ortega. Common Pitfalls Using the Normalized Compression Distance: What to Watch Out for in a Compressor. *Communications in Information and Systems*, 5(4):367–384, 2005.
- [23] M. Cebrian, M. Alfonseca, and A. Ortega. The Normalized Compression Distance is Resistant to Noise. *IEEE Transactions on Information Theory*, 53(5):1895–1900, 2007.
- [24] D. Cerra and M. Datcu. A Model Conditioned Data Compression Based Similarity Measure. In *Proceedings of the Data Compression Conference, DCC '08*, page 509, 2008.
- [25] B. Chen and Y.-T. Chen. Extractive spoken document summarization for information retrieval. *Pattern Recognition Letters*, 29(4):426–437, 2008.
- [26] X. Chen, B. Francia, M. Li, B. Mckinnon, and A. Seker. Shared Information and Program Plagiarism Detection. *IEEE Transactions on Information Theory*, 50(7):1545–1551, 2004.
- [27] Y. Chi and M. K. H. Leung. A General Shape Context Framework for Object Identification. *Computer Vision and Image Understanding*, 112(3):324–336, 2008.
- [28] R. Cilibrasi. *Statistical Inference Through Data Compression*. PhD thesis, University of Amsterdam, 2007.
- [29] R. Cilibrasi, A. L. Cruz, S. de Rooij, and M. Keijzer. CompLearn Toolkit, [Online] Available: <http://www.complearn.org/>.
- [30] R. Cilibrasi and P. Vitanyi. Clustering by Compression. *IEEE Transactions on Information Theory*, 51(4):1523–1545, 2005.
- [31] R. Cilibrasi, P. Vitanyi, and R. de Wolf. Algorithmic Clustering of Music Based on String Compression. *Computer Music Journal*, 28(4):49–67, 2004.
- [32] R. L. Cilibrasi and P. M. Vitanyi. The Google Similarity Distance. *IEEE Transactions on Knowledge and Data Engineering*, 19(3):370–383, 2007.

- [33] P. Cimiano, A. Hotho, and S. Staab. Learning Concept Hierarchies from Text Corpora using Formal Concept Analysis. *Journal of Artificial Intelligence Research*, 24(1):305–339, 2005.
- [34] J. G. Cleary and I. H. Witten. Data Compression Using Adaptive Coding and Partial String Matching. *IEEE Transactions on Communications*, 32(4):396–402, 1984.
- [35] A. Cohen, C. Bjornsson, S. Temple, G. Banker, and B. Roysam. Automatic Summarization of Changes in Biological Image Sequences Using Algorithmic Information Theory. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 31(8):1386–1403, 2009.
- [36] W. W. Cohen and Y. Singer. Context-Sensitive Learning methods for Text Categorization. *ACM Transactions on Information Systems*, 17(2):141–173, 1999.
- [37] J. R. Curran and M. Moens. Scaling Context Space. In *Proceedings of the 40th Annual Meeting on Association for Computational Linguistics*, ACL '02, pages 231–238, 2002.
- [38] M. Cuturi and J. Vert. The Context-Tree Kernel for Strings. *Neural Networks*, 18(8):1111–1123, 2005.
- [39] I. Dagan, L. Lee, and F. C. N. Pereira. Similarity-Based Models of Word Cooccurrence Probabilities. *Machine Learning*, 34(1-3):43–69, 1999.
- [40] I. Dagan, S. Marcus, and S. Markovitch. Contextual Word Similarity and Estimation from Sparse Data. In *Proceedings of the 31st annual meeting on Association for Computational Linguistics*, ACL '93, pages 164–171, 1993.
- [41] M. Dobrinkat, J. Väyrynen, T. Tapiovaara, and K. Kettunen. Normalized Compression Distance Based Measures for MetricsMATR 2010, booktitle = Proceedings of the Joint 5th Workshop on Statistical Machine Translation and MetricsMATR, year = 2010, pages = 343–348.
- [42] C. Driver and S. Clarke. An Application Framework for Mobile, Context-Aware Trails. *Pervasive and Mobile Computing*, 4(5):719–736, 2008.

- [43] A. J. Eronen, V. T. Peltonen, J. T. Tuomi, A. P. Klapuri, S. Fagerlund, T. Sorsa, G. Lorho, and J. Huopaniemi. Audio-based Context Recognition. *IEEE Transactions on Audio, Speech and Language Processing*, 14(1):321–329, 2006.
- [44] P. Ferragina, R. Giancarlo, V. Greco, G. Manzini, and G. Valiente. Compression-based Classification of Biological Sequences and Structures via the Universal Similarity Metric: Experimental Assessment. *BMC Bioinformatics*, 8, 2007.
- [45] S. Fong, D. Roussinov, and D. Skillicorn. Detecting Word Substitutions in Text. *IEEE Transactions on Knowledge and Data Engineering*, 20(8):1067–1076, 2008.
- [46] A. González-Pardo, A. Granados, D. Camacho, and F. de Borja Rodríguez. Influence of Music Representation on Compression-based Clustering. In *IEEE World Congress on Evolutionary Computation*, pages 2988 – 2995, 2010.
- [47] J. R. Goodall, W. G. Lutters, P. Rheingans, and A. Komlodi. Focusing on Context in Network Traffic Analysis. *IEEE Computer Graphics and Applications*, 26(2):72–80, 2006.
- [48] A. Granados, D. Camacho, and F. de Borja Rodríguez. Is the contextual information relevant in text clustering by compression? *Expert Syst. Appl. (Submitted February 2011)*.
- [49] A. Granados, M. Cebrián, D. Camacho, and F. de Borja Rodríguez. Reducing the Loss of Information through Annealing Text Distortion. *IEEE Transactions on Knowledge and Data Engineering*, 23(7):1090–1102, July 2011.
- [50] A. Granados, M. Cebrián, D. Camacho, and F. B. Rodríguez. Evaluating the Impact of Information Distortion on Normalized Compression Distance. In *Proceedings of the 2nd international Castle meeting on Coding Theory and Applications*, ICMCTA '08, Lecture Notes in Computer Science, pages 69–79, 2008.
- [51] A. Granados, R. Martínez, D. Camacho, and F. de Borja Rodríguez. Relevance of Contextual Information in Compression-Based Text Clustering. In *Proceedings of the 11th international conference on Intelligent Data Engineering and Automated Learning*, IDEAL '10, Lecture Notes in Computer Science, pages 259–266, 2010.

- [52] A. Granados, R. Martínez, D. Camacho, and F. Rodríguez. Improving the Accuracy of the Normalized Compression Distance combining Document Segmentation and Document Distortion. *IEEE Transactions on Knowledge and Data Engineering* (Submitted October 2011).
- [53] Z. Guodong. Modeling of Long Distance Context Dependency. In *Proceedings of the 20th international conference on computational linguistics*, COLING '04, page 92, 2004.
- [54] M. A. Hearst and C. Plaunt. Subtopic structuring for full-length document access. In *Proceedings of the 16th annual international ACM SIGIR conference on Research and development in information retrieval*, SIGIR '93, pages 59–68, 1993.
- [55] R. M. Hegde, J. Kurniawan, and B. D. Rao. On the Design and Prototype Implementation of a Multimodal Situation Aware System. *IEEE Transactions on Multimedia*, 11(4):645–657, 2009.
- [56] S. Helmer. Measuring the structural similarity of semistructured documents using entropy. In *Proceedings of the 33rd international conference on Very large data bases*, pages 1022–1032, 2007.
- [57] M. Hu, A. Sun, and E.-P. Lim. Comments-oriented Document Summarization: Understanding Documents with Readers' Feedback. In *Proceedings of the 31st annual international ACM SIGIR conference on Research and development in information retrieval*, SIGIR '08, pages 291–298, 2008.
- [58] C.-L. Huang and C.-H. Wu. Generation of Phonetic Units for Mixed-Language Speech Recognition Based on Acoustic and Contextual Analysis. *IEEE Transactions on Computers*, 56(9):1225–1233, 2007.
- [59] D. A. Huffman. A Method for the Construction of Minimum-Redundancy Codes. *Proceedings of the Institute of Radio Engineers*, 40(9):1098–1101, 1952.
- [60] IMDB, Internet Movie Database, [Online] Available: <http://www.imdb.com/>.
- [61] T. Joachims. Text Categorization with Support Vector Machines: Learning with Many Relevant Features. In *Proceedings of the 10th European Conference on Machine Learning*, pages 137–142, 1998.

- [62] M. Kaszkiel and J. Zobel. Passage Retrieval Revisited. In *Proceedings of the 20th annual international ACM SIGIR conference on Research and development in information retrieval*, SIGIR '97, pages 178–185, 1997.
- [63] D. Knaus, E. Mittendorf, P. Schäuble, and P. Sheridan. Highlighting relevant passages for users of the interactive spider retrieval system. In *TREC-4 Proceedings*, 1996.
- [64] Y. Ko and J. Seo. An Effective Sentence-Extraction Technique using Contextual Information and Statistical Approaches for Text Summarization. *Pattern Recognition Letters*, 29(9):1366–1371, 2008.
- [65] A. Kocsor, A. Kertész-Farkas, L. Kaján, and S. Pongor. Application of Compression-based Distance Measures to Protein Sequence Classification: a Methodological Study. *Bioinformatics*, 22(4):407–412, 2006.
- [66] A. Kolmogorov. Three Approaches to the Quantitative Definition of Information. *Problems Information Transmission*, 1(1):1–7, 1965.
- [67] H. Kozima. Text Segmentation Based on Similarity between Words. In *Proceedings of the 31st annual meeting on Association for Computational Linguistics*, ACL '93, pages 286–288, 1993.
- [68] A. Kraskov, H. Stoegbauer, R. Andrzejak, and P. Grassberger. Hierarchical Clustering using Mutual Information. *Europhysics Letters*, 70(2):278–284, 2005.
- [69] N. Krasnogor and D. Pelta. Measuring the Similarity of Protein Structures by means of the Universal Similarity Metric. *Bioinformatics*, 20(7):1015–1021, 2004.
- [70] K. L. Kwok, L. Grunfeld, and D. D. Lewis. TREC-3 Ad-Hoc, Routing Retrieval and Thresholding Experiments using PIRCS. In *Proceedings of TREC-3*, pages 247–255, 1995.
- [71] O. Kwon and J. Kim. Concept Lattices for Visualizing and Generating User Profiles for Context-Aware Service Recommendations. *Expert Systems with Applications*, 36(2):1893–1902, 2009.
- [72] M. Lan, C. Tan, J. Su, and Y. Lu. Supervised and Traditional Term Weighting Methods for Automatic Text Categorization. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 31(4):721–735, 2009.

- [73] R. Y. K. Lau, P. D. Bruza, and D. Song. Towards a Belief-Revision-based Adaptive and Context-Sensitive Information Retrieval System. *ACM Transactions on Information Systems*, 26(2):1–38, 2008.
- [74] K.-F. Lee. Context-dependent Phonetic Hidden Markov models for Speaker-Independent Continuous Speech Recognition. In *Readings in speech recognition*, pages 347–366, 1990.
- [75] M. Li, X. Chen, X. Li, B. Ma, and P. Vitanyi. The Similarity Metric. *IEEE Transactions on Information Theory*, 50(12):3250–3264, 2004.
- [76] M. Li and P. Vitanyi. *An introduction to Kolmogorov complexity and its applications*. Springer-Verlag, 3rd edition, 1997.
- [77] R.-L. Liu. Context-Based Term Frequency Assessment for Text Classification. In *Proceedings of the 10th Pacific Rim International Conference on Artificial Intelligence*, PRICAI '08, pages 1004–1009, 2008.
- [78] R.-L. Liu and Y.-L. Lu. Incremental Context Mining for Adaptive Document Classification. In *Proceedings of the eighth ACM SIGKDD international conference on Knowledge discovery and data mining*, KDD '02, pages 599–604, 2002.
- [79] H. P. Luhn. The Automatic Creation of Literature Abstracts. *IBM Journal of Research and Development*, 2:159–165, 1958.
- [80] Z. Ma, G. Pant, and O. R. L. Sheng. Interest-based Personalized Search. *ACM Transactions on Information Systems*, 25(1):5, 2007.
- [81] G. N. N. Martin. Range Encoding: an Algorithm for Removing Redundancy from a Digitised Message. In *Video and Data Recording Conference*, 1979.
- [82] R. Martínez, M. Cebrián, F. de Borja Rodríguez, and D. Camacho. Contextual Information Retrieval based on Algorithmic Information Theory and Statistical Outlier Detection. In *IEEE Information Theory Workshop*, 2008.
- [83] R. M. Martínez. An introduction to content-based Information Retrieval by Normalized Compression Distance. Technical report, Universidad Autónoma de Madrid, 2009.
- [84] MedlinePlus Health Information, MedlinePlus website, U.S. National Library of Medicine and National Institutes of Health. [Online] Webpage: <http://medlineplus.gov/>.

- [85] J. L. Melville, J. F. Riley, and J. D. Hirst. Similarity by Compression. *Journal of Chemical Information and Modeling*, 47(1):25–33, 2007.
- [86] E. Mittendorf and P. Schäubel. Document and Passage Retrieval based on Hidden Markov Models. In *Proceedings of the 17th annual international ACM SIGIR conference on Research and development in information retrieval*, SIGIR '94, pages 318–327, 1994.
- [87] A. Moffat. Implementing the PPM Data Compression Scheme. *IEEE Transactions on Communications*, 38(11):1917–1921, 1990.
- [88] R. J. Mooney and R. Bunescu. Mining Knowledge from Text using Information Extraction. *ACM SIGKDD Explorations Newsletter - Natural Language Processing and Text Mining*, 7(1):3–10, 2005.
- [89] G. Mori, S. Belongie, and J. Malik. Efficient Shape Matching Using Shape Contexts. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 27(11):1832–1837, 2005.
- [90] G. Mori and J. Malik. Recovering 3D Human Body Configurations Using Shape Contexts. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 28(7):1052–1062, 2006.
- [91] M. Nykter, N. D. Price, M. Aldana, S. A. Ramsey, S. A. Kauffman, L. E. Hood, O. Yli-Harja, and I. Shmulevich. Gene Expression Dynamics in the Macrophage Exhibit Criticality. *Proceedings of the National Academy of Sciences*, 105(6):1897–1900, 2008.
- [92] J. J. Odell. *The Use of Context in Large Vocabulary Speech Recognition*. PhD thesis, University of Cambridge, 1995.
- [93] [Online] <http://www.opensubtitles.org/en>. *The biggest open subtitles database*.
- [94] D. Ornstein and B. Weiss. Entropy and data compression schemes. *IEEE Transactions on Information Theory*, 39(1):78–83, 1993.
- [95] B. Pang and L. Lee. Opinion Mining and Sentiment Analysis. *Foundations and Trends in Information Retrieval*, 2(1-2):1–135, 2008.
- [96] J. Pascoe. Adding Generic Contextual Capabilities to Wearable Computers. In *Proceedings of the 2nd IEEE International Symposium on Wearable Computers*, ISWC '98, page 92, 1998.

- [97] I. Pavlov. LZMAX, [Online] Available: <http://www.7-zip.org/sdk.html>.
- [98] X. Quan, L. Wenyin, and B. Qiu. Term Weighting Schemes for Question Categorization. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 33(5):1009–1021, 2011.
- [99] D. Ravichandran and E. Hovy. Learning Surface Text Patterns for a Question Answering System. *Proceedings of the 40th Annual Meeting on Association for Computational Linguistics*, pages 41–47, 2001.
- [100] B. Rhodes. Using Physical Context for Just-in-Time Information Retrieval. *IEEE Transactions on Computers*, 52(8):1011–1014, 2003.
- [101] L. Rocha, a. Fernando Mour A. Pereira, M. A. Gonçalves, and W. M. Jr. Exploiting Temporal Contexts in Text Classification. In *Proceeding of the 17th ACM Conference on Information and Knowledge Management*, CIKM '0, pages 243–252, 2008.
- [102] B. Y. Ryabko. Data Compression by means of a 'Book Stack'. *Problems of Information Transmission*, 16:265–269, 1980.
- [103] D. Salomon. *Data Compression: The Complete Reference*. Springer, 2004.
- [104] G. Salton. *Automatic Text Processing: the Transformation, Analysis, and Retrieval of Information by Computer*. Addison-Wesley Longman Publishing Co., 1989.
- [105] G. Salton, J. Allan, and C. Buckley. Approaches to Passage Retrieval in Full Text Information Systems. In *Proceedings of the 16th annual international ACM SIGIR conference on Research and development in information retrieval*, SIGIR '93, pages 49–58, 1993.
- [106] G. Salton, A. Wong, and C.-S. Yang. A Vector Space Model for Automatic Indexing. *Communications of the ACM*, 18(11):613–620, 1975.
- [107] C. C. Santos, J. Bernardes, P. M. B. Vitanyi, and L. Antunes. Clustering Fetal Heart Rate Tracings by Compression. In *Proceedings of the 19th IEEE Symposium on Computer-Based Medical Systems*, CBMS '06, pages 685–690, 2006.
- [108] B. Schilit, N. Adams, and R. Want. Context-Aware Computing Applications. In *Proceedings of the 1994 First Workshop on Mobile Computing Systems and Applications*, WMCSA '94, pages 85–90, 1994.

- [109] W. T. Scott. A new approach to data mining for software design. URL citeseer.ist.psu.edu/727888.html.
- [110] F. Sebastiani. Machine learning in automated text categorization. *ACM Computing Surveys*, 34(1):1–47, 2002.
- [111] J. Seward. BZIP2, [Online] Available: <http://bzip.org/>.
- [112] C. Shah and W. B. Croft. Evaluating High Accuracy Retrieval Techniques. In *Proceedings of the 27th annual international ACM SIGIR conference on Research and development in information retrieval*, SIGIR '04, pages 2–9, 2004.
- [113] C. E. Shannon. A Mathematical Theory of Communication. *The Bell System Technical Journal*, 27:379–423, 623–656, 1948.
- [114] O. Shiyan, C. S.-G. Khoo, and D. H. Goh. Design and Development of a Concept-based Multi-Document Summarization System for Research Abstracts. *Journal of Information Science*, 34(3):308–326, 2008.
- [115] A. Sieg, B. Mobasher, and R. Burke. Web Search Personalization with Ontological User Profiles. In *Proceedings of the sixteenth ACM conference on Conference on information and knowledge management*, CIKM '07, pages 525–534, 2007.
- [116] C. Silva and B. Ribeiro. The Importance of Stop Word Removal on Recall Values in Text Categorization. In *Proceedings of the International Joint Conference on Neural Networks*, 2003.
- [117] M. Speretta and S. Gauch. Personalized Search Based on User Search Histories. In *Proceedings of the 2005 IEEE/WIC/ACM International Conference on Web Intelligence*, WI '05, pages 622–628, 2005.
- [118] J.-H. Su, H.-H. Yeh, P. S. Yu, and V. S. Tseng. Music Recommendation Using Content and Context Information Mining. *IEEE Intelligent Systems*, 25:16–26, 2010.
- [119] N. Svengård and P. Nordin. Automated Aesthetic Selection of Evolutionary Art by Distance Based Classification of Genomes and Phenomes Using the Universal Similarity Metric. In *EvoWorkshops*, pages 447–456, 2004.
- [120] L. Tamine-Lechani, M. Boughanem, and M. Daoud. Evaluation of Contextual Information Retrieval Effectiveness: Overview of Issues and Research. *Knowledge and Information Systems*, 24(1):1–34, 2010.

- [121] G. Telles, R. Minghim, and F. Paulovich. Normalized Compression Distance for Visual Analysis of Document Collections. *Computers & Graphics*, 31(3):327–337, 2007.
- [122] J. Tiedemann and J. Mur. Simple is Best: Experiments with Different Document Segmentation Strategies for Passage Retrieval. In *Proceedings of the 2nd workshop on Information Retrieval for Question Answering*, pages 17–25, 2008.
- [123] J. T. T.M. Cover. *Elements of Information Theory*. Wiley and Sons, 1991.
- [124] A. Turing. On Computable Numbers, with an Application to the Entscheidungsproblem. *Proceedings of the London Mathematical Society*, 2(42):230–265, 1936.
- [125] UCI Knowledge Discovery in Databases Archive, Information and Computer Science, University of California, Irvine. [Online] Available: <http://kdd.ics.uci.edu/>.
- [126] C. Van Rijsbergen. *Information Retrieval*. Butterworth-Heinemann Newton, MA, USA, 1979.
- [127] R. Varadarajan and V. Hristidis. A System for Query-Specific Document Summarization. In *Proceedings of the 15th ACM international conference on Information and knowledge management, CIKM '06*, pages 622–631, 2006.
- [128] S. Verdú and T. Weissman. The Information Lost in Erasures. *IEEE Transactions on Information Theory*, 54(11):5030–5058, 2008.
- [129] A. Vinciarelli. Noisy Text Categorization. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 27(12):1882–1895, 2005.
- [130] J. Weeds and D. Weir. Co-occurrence Retrieval: A Flexible Framework for Lexical Distributional Similarity. *Computational Linguistics*, 31(4):439–475, 2005.
- [131] S. Wehner. Analyzing Worms and Network Traffic using Compression. *Journal of Computer Security*, 15:303–320, 2007.
- [132] S.-S. Weng, B. Lin, and W.-T. Chen. Using contextual information and multidimensional approach for recommendation. *Expert Systems with Applications*, 36(2):1268–1279, 2009.

- [133] W. Wilbur and K. Sirotkin. The Automatic Identification of Stop Words. *Journal of Information Science*, 18(1):45, 1992.
- [134] R. Wilkinson. Effective Retrieval of Structured Documents. In *Proceedings of the 17th annual international ACM SIGIR conference on Research and development in information retrieval*, SIGIR '04, pages 311–317, 1994.
- [135] W. Xi, R. Xu-Rong, C. S. Khoo, and E.-P. Lim. Incorporating Window-based Passage-Level Evidence in Document Retrieval. *Journal of Information Science*, 27(2):73–80, 2001.
- [136] H. Xiong, G. Pandey, M. Steinbach, and V. Kumar. Enhancing Data Analysis with Noise Removal. *IEEE Transactions on Knowledge and Data Engineering*, 18(3):304–319, 2006.
- [137] Y. Yang. Noise reduction in a statistical approach to text categorization. In *Proceedings of the 18th annual international ACM SIGIR conference on Research and development in information retrieval*, SIGIR '95, pages 256–263, 1995.
- [138] Y. Yang and J. Wilbur. Using Corpus Statistics to Remove Redundant Words in Text Categorization. *Journal of the American Society for Information Science*, 47(5):357–369, 1996.
- [139] X. Zhang, Y. Hao, X. Zhu, and M. Li. Information Distance from a Question to an Answer. In *Proceedings of the 13th ACM SIGKDD international conference on Knowledge discovery and data mining*, KDD '07, pages 874–883, 2007.
- [140] X. Zhang, Y. Hao, X.-Y. Zhu, and M. Li. New Information Distance Measure and its Application in Question Answering System. *Journal of Computer Science and Technology*, 23(4):557–572, 2008.
- [141] G. K. Zipf. *The Psychobiology of Language*. Houghton-Mifflin, 1935.
- [142] G. K. Zipf. *Human Behavior and the Principle of Least Effort*. Addison-Wesley (Reading MA), 1949.
- [143] J. Ziv and A. Lempel. Compression of Individual Sequences by Variable Rate Coding. *IEEE Transactions on Information Theory*, 24(5):530–536, 1978.

- [144] J. Ziv and A. Lempel. A Universal Algorithm for Sequential Data Compression. *IEEE Transactions on Information Theory*, 23(3):337–343, 2003.
- [145] J. Zobel, A. Moffat, R. Wilkinson, and R. Sacks-Davis. Efficient Retrieval of Partial Documents. *Information Processing and Management*, 31:361–377, 1995.